Using Generative AI for Advocacy: Risks and Benefits

Summary

Generative artificial intelligence (AI) refers to systems that can create human-like text, images, and other content in response to prompts. This document will explain how to responsibly use tools like ChatGPT, Claude, Gemini, Copilot, and Meta AI to assist with writing, research, social media posts, strategy development, and campaign planning.

Generative AI is a potentially transformative technology that can dramatically increase efficiency and expand capacity. At the same time, the hazards are clear: Loose privacy standards, enormous energy needs, fake – but believable – content creation, “hallucinations,” and more make AI a challenging technology to use ethically. However, AI is here to stay, and it is quickly inserting itself into every aspect of technological life. To keep up with your adversaries and allies, it is necessary to learn to use AI responsibly while implementing strong guardrails for privacy protection and establishing strict organizational policies around data sharing, content quality, and human oversight.

Benefits of Generative AI for Advocacy

Content Creation and Communication
Generative AI excels at quick content development, enabling users to quickly create draft content including letters to policymakers, social media posts, talking points for spokespeople, internal memos and, depending on the platform, images. While not without risks, using AI to create drafts of such content can be a reasonable starting point for small organizations or groups, especially those that lack staff, time, or financial resources.

Analysis from Multiple Sources
AI tools synthesize information from multiple source documents to create comprehensive outlines and work plans. Some platforms allow users to create folders or projects that store source documents, creating a knowledgebase to draw from when creating outlines or draft content.

This capability is particularly valuable when incorporating research from various reports, studies, and policy documents; or when writing complex policy documents that might otherwise take hours to analyze.

This work to synthesize and write content can be *assisted* by AI but should not be *accomplished* by AI, the latter of which risks publishing inaccurate or untrusted information in an inauthentic voice.

Audience-Specific Messaging
Organizations can use AI to quickly draft tailored messaging for different audiences and demographics using templates. This allows for more personalized communication without requiring extensive additional staff time for customization – though it is of course essential for human staff to edit and review any content drafted by AI.

How AI Projects and Knowledge Bases Work

About “Projects”
AI tools don't remember previous conversations, so each interaction starts from scratch. Projects or folders allow you to create a persistent knowledge base that the AI can reference across multiple conversations. When you upload documents to a project, the AI can draw connections between different sources and provide more comprehensive, contextually-aware responses.

How AI "Learns" from Your Documents
AI doesn't actually learn or update from your uploads - instead, it temporarily incorporates your documents into its working memory for that session. Think of it like giving the AI a stack of reference materials to consult while answering questions. The AI will search through your uploaded documents to find relevant information and combine it with its general training knowledge.

AI tools can only work with what you explicitly provide. If you upload incomplete information or documents that contradict each other, the AI may give you incomplete or conflicting responses. It won't know what you haven't told it.

AI Tools Are Experimental and Error-Prone
These tools are new, rapidly evolving, and make frequent mistakes. They can:

- Confidently state false information that sounds plausible
- Miss key details in documents you've uploaded
- Combine information incorrectly from multiple sources
- Generate outdated information based on their training data
- Fail to understand nuanced political or cultural contexts

The AI Doesn't Always Answer Your Actual Question
AI tools often respond to what they think you're asking rather than what you actually asked. They may:

- Focus on the wrong aspect of a complex question
- Provide generic advice when you need specific guidance
- Miss the strategic or political implications of your situation
- Give you technically correct but practically useless information

Always Verify and Cross-Check

- Fact-check all claims, statistics, and citations
- Test AI recommendations on a small scale before full implementation

- Consult human experts for strategic decisions
- Review AI output with fresh eyes - does it actually address your needs?

How to Communicate Effectively with AI Tools

Be Specific and Direct
Instead of: "Help with our campaign" Try: "Draft three key talking points about our opposition to the proposed energy project, focusing on environmental impact, using a tone appropriate for voters in key districts"

Provide Context

- Explain who your audience is
- Specify the format you need (email, social post, policy brief)
- Include relevant constraints (word count, deadline, tone)
- Mention what you've already tried or what hasn't worked

Use Follow-Up Questions

- "Can you make this more specific to [your location/issue]?"
- "This seems too generic - can you make it more compelling?"
- "What evidence would make this argument stronger?"

Ask for Citations and Sources
Always include: "Please cite your sources and provide links where possible" or "What evidence supports this recommendation?"

Iterate and Refine
Don't expect perfect results on the first try. Use responses like:

- "This is close, but can you focus more on [specific aspect]?"
- "The tone isn't quite right - can you make it more [urgent/professional/accessible]?"
- "Can you provide three different approaches to this same message?"

Test Your Prompts
Ask the same question in different ways to see if you get consistent, useful responses. If results vary wildly, your question may need to be more specific.

Use Cases

Building Knowledge Bases with Projects/Folders

Setting Up Organizational Knowledge

- In the paid versions of ChatGPT or Claude, you can create “Projects” for each campaign, issue, or project you are working on, where you can upload background research, policy documents, approved messaging, and other resources. These resources act as a memory and knowledge base for your work.
- Upload public documents only: published reports, public statements, press releases, and general research materials
- Use these knowledge bases to generate consistent talking points, background summaries, project outlines, and educational content
- Example: A folder containing public reports, published studies, your organization's past public statements, and other non-sensitive materials can help generate draft content and outlines

Best Practices

- Rule of thumb: Only upload materials you'd be comfortable making public
- Regularly audit folder contents to remove outdated information
- Create separate projects for different spheres of work (policy work, communications, organizational strategy)

Safe Data Practices - Working Without Sensitive Information

Document Sanitization Workflow:

- Create "clean" versions of documents with sensitive information removed (names, addresses, internal strategy details, donor information)
- Use placeholder text like "[ORGANIZATION NAME]" or "[LOCAL REPRESENTATIVE]" that you can customize after AI generates content
- When possible, work with publicly available data and research rather than internal documents, and only use internal documents that you’d be comfortable making public

Template-Based Approach

- Develop messaging templates that AI can customize for different contexts
- Example: "Create three versions of this message - one for suburban voters, one for rural communities, and one for urban areas"
- Use AI to adapt public-facing content rather than create from scratch

Research and Analysis

Multi-Source Synthesis:

- Upload multiple public documents on the same issue and ask AI to summarize, identify common themes, highlight conflicting viewpoints, or note gaps in coverage
- Generate literature reviews from publicly available academic papers and policy reports
- Create comparative analyses of policy papers using publicly available documents

Opposition Research:

- Analyze publicly available voting records, statements, and policy positions
- Generate fact-check summaries of public claims by opponents
- Always verify AI findings with original sources

Content Adaptation and Translation

Audience-Specific Messaging:

- Take existing public content and adapt it for different platforms (social media posts, newsletters, policy briefs)
- Adjust reading level and tone for different audiences
- Generate multiple versions for A/B testing

Multilingual Outreach:

- Translate public-facing materials (with human verification)
- Adapt cultural messaging for different communities
- Create culturally appropriate versions of the same core message

Risks and Challenges

Data and Privacy Concerns

Sensitive Information Exposure
Organizations must assume that any information shared with AI tools will become public.

The most critical risk involves uploading sensitive data. Names, addresses, contact information, and other personally identifiable information could potentially be exposed through platform tracking of uploaded data or via data breaches, or it could be incorporated into AI training datasets. Which means this information can be used inappropriately by anyone using the AI platform.

Strategic Intelligence Leaks
Campaign strategies, lists of names or email addresses, internal planning documents, and sensitive details about advocacy efforts should never be uploaded to AI platforms. The ways in which AI platforms digest, use, and share uploaded data are opaque at best, meaning there’s no clarity on whether or how uploaded data maybe be displayed to other users. Competitors or opposition groups could potentially gain access to this information, compromising campaign effectiveness.

Legal Vulnerabilities
AI-generated content may infringe on copyright, create misleading representations, or expose organizations to legal liability. As always, humans **must **review and edit such content before sharing it publicly. The legal landscape around AI-generated content remains unsettled, creating additional uncertainty.

Content Quality and Authenticity Issues

Generic "AI-Speak"
AI tools, attempting to mimic authentic speech, can often end up using patterns and phrasing that can feel impersonal or inauthentic. Organizations must be sure to customize AI-generated drafts to ensure their work is in their own authentic voice.

Factual Inaccuracies
AI systems can confidently present false information, generate non-existent statistics, or misrepresent facts—a phenomenon known as "hallucination." All AI-generated content requires careful fact-checking and verification.

Bias Amplification
AI systems can perpetuate and amplify harmful stereotypes present in their training data, subtly inserting biased language or assumptions into messaging that could contradict organizational values and may go unnoticed without careful review.

Loss of Authentic Voices
Too much reliance on AI-generated content risks replacing genuine voices with artificial alternatives, potentially undermining the authenticity that gives your work its moral authority and emotional resonance.

Limitations with Numerical Data and Calculations
While AI tools can be helpful for analyzing trends and patterns in data, they have significant limitations when working with precise numerical calculations and datasets.

AI systems may introduce errors in mathematical computations, misinterpret numerical relationships, or provide results that appear accurate but contain subtle mistakes. Users should always verify numerical outputs through independent calculations or specialized analytical software, especially for data that will inform important decisions.

When using AI for data analysis, treat the results as a starting point for further verification rather than as definitive findings. For critical numerical work, consider using dedicated statistical software or spreadsheet applications alongside AI tools rather than relying on AI alone.

Operational and Strategic Risks

Diminished Human Judgment
Too much dependence on AI tools can lead to a reduced focus on critical thinking and strategic analysis; it’s easy to just let the robots do the work if you’re not careful! Empathy, critical thought, and relationship-building remain essential to effective advocacy and there are no replacements for them.

Financial Considerations
Premium AI subscriptions – often necessary for anything beyond one-off questions – can incur significant ongoing costs, which may be particularly challenging for smaller organizations to bear.

Best Practices for Safe AI Adoption

*Any AI-created materials **must** undergo a process of human review before publication!*

Research
To avoid an AI tool sharing false information with you while doing research, in your prompt, always tell it to cite its findings for you. Verify that all links it shares are accurate before using the citation. For example, make sure the link the AI references actually works and leads you to where the AI found the research.

Data Protection Protocols
Establish clear guidelines for what information can and cannot be shared with AI tools. Create workflows that strip sensitive data before uploading and maintain secure processes for handling confidential information.

Human Oversight Requirements
Ensure that a human reviews – and edits, if necessary – all AI-assisted content before publication or distribution. Designate staff members responsible for fact-checking and ensuring accuracy of AI-assisted work.

Policies
Consider developing an AI policy that is transparently shared on your website. It can include commitments to mark all AI-assisted content with a label or footnote.

Staff Training and Education
Provide comprehensive training on AI capabilities, limitations, and risks. Ensure staff understand both the potential and the pitfalls of these tools.

System Settings and Privacy Configurations
Most platforms offer various privacy and data usage settings that users should review carefully.

Common settings include options to "improve the model for everyone" (which may allow your conversations to be used in training data), memory features that save information across sessions, and data retention policies. For maximum privacy protection, users should generally disable data sharing for model improvement and turn off persistent memory features.

However, even with privacy settings enabled, users should still avoid uploading confidential data or sensitive organizational information.

Staying Current
The AI landscape evolves rapidly, making it essential for organizations to stay informed about new capabilities and emerging risks. Reliable sources for ongoing AI security and capability updates include the National Institute of Standards and Technology (NIST), though their resources may be affected by recent budget cuts. Academic institutions such as Stanford's Human-Centered AI Institute and MIT's Computer Science and Artificial Intelligence Laboratory publish research on AI safety and capabilities.

Cost
The adage “you get what you pay for” is particularly true with AI. The free versions of ChatGPT, Claude, Gemini and others all come with significant usage limitations, thanks to resource-intensive nature of the systems. All of these platforms offer paid tiers at around $20/month, which grant more time and extra features like folders/projects and, in some cases, the ability to opt out of having your chats used to train AI models.

While this monthly fee will feel steep to many, if you come to rely on AI tools for your daily work it is likely worth the cost.

Privacy-Focused Alternatives
Several platforms offer access to popular AI models through privacy-focused interfaces. Services like Duck.ai and Kagi's AI features provide access to leading language models while implementing stronger privacy safeguards, such as not storing conversation histories or using interactions for model training.

However, users should still exercise the same caution regarding sensitive data, as queries are still be processed by the underlying AI models (such as ChatGPT, Claude, Llama, and Google’s Gemini) even if the intermediary platform doesn't retain the information.

Revision #7
Created 30 October 2025 20:13:50 by Josh
Updated 30 October 2025 22:06:08 by Josh