MEMORANDUM FOR: All State of Maryland Executive Branch Staff
FROM: Katie Savage, Secretary of Department of Information Technology, Maryland Chief Information Officer, and Chair of the Governor's AI Subcabinet
SUBJECT: **Interim Guidance for the Responsible Use of Commercial Generative Artificial Intelligence Tools**
VERSION: 1.0
EFFECTIVE DATE: 5/9/24
CHANGE LOG: N/A

I. Purpose & background

Generative Artificial Intelligence (GenAI) refers to a set of AI technologies and models that learn patterns and relationships from massive datasets to generate novel content (text, images, audio, video, and code) based on user prompts. GenAI holds tremendous promise, and if properly harnessed, can help the State of Maryland's workforce better serve residents and take drudgery out of workflows. At the same time, there are significant privacy, security, bias, accuracy, and legal risks to take into account.

- Governance-by-guidelines, in service of experimentation and learning, relies on careful human judgment to balance productivity gains with risks.
- Given the technology's rate of change, and in order to enable use and experimentation while centering responsible practices, State government employees must follow these interim guidelines if they leverage reputed commercial, standalone GenAI solutions like OpenAI's ChatGPT or DALL-E, Google's Gemini, Anthropic's Claude, or Microsoft's Copilot, whether free versions or paid licenses.
- For the time being, while purchasing of any paid, premium, or enterprise versions of these tools is delegated to the discretion of individual agencies, it must still go through the existing DoIT intake process.
- Individual agencies may choose to add additional guidelines or restrictions as necessary given their particular context, but should first consult the AI Subcabinet (email us at AI@maryland.gov), which is tasked by the Governor's AI Executive Order to ensure the State is well positioned to responsibly, ethically, and productively procure, leverage, govern, and deploy AI.
- We anticipate this interim guidance to be incorporated into a more comprehensive Responsible Use policy for AI (including but not limited to GenAI). That policy will be developed and shared later in 2024 upon further study by DoIT and the AI Subcabinet.
- Customized software or services through an IT procurement process are not covered by this document.
- These guidelines will be periodically reviewed and updated as the technologies and products evolve and opportunities and risks change.
- In time, an eventual enterprise AI adoption strategy, along with a more formally developed approach to building an "allow list", will further specify what tools are available, and not available, for state employee use.

II. Direction

In order to use reputed, commercially available GenAI tools, each State employee must understand their accountability for responsible GenAI use and will be required to read carefully through this document. We will ask you to complete intermittent surveys to share more about your use of GenAI tools, as an input into where to invest in the future on tooling and governance. We also anticipate beginning to offer basic GenAI training modules to supplement these guidelines in mid-2024.

III. AI principles

All use of AI - including Generative AI - must adhere to the following principles. Further direction to adhere to these principles will be developed by the AI Subcabinet and DoIT. In the interim, all State employees should use their professional and discretionary judgment as it relates to the guidelines in this document. Furthermore, it is your responsibility to think critically about use of these emerging technologies; if you have concerns, be prudent and err on the side of caution.

1. **Fairness and Equity**. The State's use of AI must take into account the fact that AI systems can perpetuate harmful biases, and take steps to mitigate those risks, in order to avoid discrimination or disparate impact to individuals or communities based on their race, color, ethnicity, sex, religion, age, ancestry or national origin, disability, veteran status, marital status, sexual orientation, gender identity, genetic information, or any other classification protected by law.
2. **Innovation**. When used responsibly and in human-centered and mission-aligned ways, AI has the potential to be a tremendous force for good. The State commits to exploring ways AI can be leveraged to improve State services and resident outcomes.
3. **Privacy**. Individuals' privacy rights should be preserved by design in the State's use of AI, while ensuring that data creation, collection, and processing are secure and in line with all applicable laws and regulations.
4. **Safety, security, and resiliency**. AI presents new challenges and opportunities for ensuring the safety and security of Maryland residents, infrastructure, systems, and data. The State commits to adopting best practice guidelines and standards to surface and mitigate safety risks stemming from AI, while ensuring AI tools are resilient to threats.
5. **Validity and reliability**. AI systems can change over time. There should be mechanisms in place to ensure that these systems are working as intended, together with accurate outputs and robust performance.
6. **Transparency, accountability, and explainability**. The State's use of AI should be clearly and regularly documented and disclosed, in order to enable accountability. The outputs of AI systems in use by the State should be explainable and interpretable to oversight bodies and residents, with clear human oversight.

IV. Guidelines

1. **PI & Sensitive information:** Do not input personal information (PI)[1] or protected health information (PHI)[2], or any information about individuals, into any commercial GenAI tool - even if it's "anonymized" or "de-identified." In addition to PI and PHI, do not input confidential or sensitive content. Many GenAI tools allow you to upload documents as reference information that the tool will leverage to respond to your prompts. You are only permitted to upload documents that are already publicly available, e.g., any content that you have a reasonable basis to believe is lawfully made available to the general public from federal, state, or local government records; widely distributed media; or accessible online without the need for special qualifications, permissions, or privileges. Do not upload non-public communications, protected/internal-only documents, attorney work product, pre-decisional/deliberative documents, documents that contain PI/PHI, or emails or chats with colleagues. Immediately speak with your agency's Privacy Officer if you believe PI may have been or was actually leaked, or if you have any questions about what constitutes PI, PHI, or sensitive content.
2. **Decision-making and evaluations:** Staff are not permitted to use commercial, off-the-shelf GenAI tools to make sensitive decisions or evaluations concerning individual benefits, credentialing, vetting, hiring, legal, or civil investigations or enforcement related actions, or any State activities affecting individual rights or safety.
3. **Fact-check:** GenAI can create content that seems real but is not. Many GenAI systems are also only trained on data up to a certain date, so questions asking about current events may not yield relevant results. Ensure that you fact-check output using a different source prior to using the outputs in your work product. This is particularly important anywhere accuracy is paramount or outputs will be used publicly. Keep in mind that you are responsible for the veracity of your work and any inaccuracies therein, even if you leverage GenAI tools to help in its production.
4. **Review for bias:** AI tools can reflect existing societal biases captured from their training data. That may lead to tools outputting assumptions based on past stereotypes that need to be corrected. An example is generating gender-biased job descriptions given traditional stereotypes associated with certain roles; or, when creating interview shortlists from uploaded resumes, scoring resumes with certain names more highly than others, even with all other things being equal. Be aware of individual biases and ask questions of an AI tool in a neutral, open-minded manner.
5. **Transcription tools:** Do not use GenAI tools to transcribe or summarize meetings where sensitive topics or PI are discussed. You may find that some transcription tools are blocked. If you do use an AI transcription tool, ensure you've documented in the event invite that the meeting will be transcribed, state as much at the beginning of the meeting, and disable transcription if any attendee does not consent. Public meetings and hearings are good places to leverage transcription tools, so as to easily make available such content.

[1] PI may include information about residents, colleagues, or yourself, and is information that can be used to distinguish or trace someone's identity, either by itself, or if combined with other information. You can learn more about what constitutes PI under the MD Protection of Information in Government Act.
[2] More detail on what constitutes PHI is available under the HIPAA Privacy Rule.

6. **Access and accounts:** Create accounts on these tools using your @maryland.gov email and use them only for work-related purposes, separate from any personal use of these tools. Select options in tools that limit data retention and opt out of your input information being used to further train models. All use of conditionally approved tools will be on the web, not through downloaded desktop or mobile apps. While offboarding from State service, or when you no longer need the GenAI service, you must submit a request to the GenAI tool to delete your account.
7. **Terms and conditions (T&Cs):** Read through the T&Cs you agree to before you use GenAI tools. The state does not currently have agreements in place for common GenAI systems, like ChatGPT or Google Gemini or Microsoft Copilot. If you choose to use GenAI for state work and agree to the T&Cs without a State agreement in place, you are responsible for complying with those T&Cs – in addition to these guidelines. Individual agencies may impose additional conditions or limitations of use based on their unique context and the tools' T&Cs.
8. **Disclosures and citation:** To build trust in State government and to adhere to the transparency principle, you must disclose and cite GenAI use in certain circumstances. We will evolve the form and context of these disclosures over time as capabilities change, and to cover GenAI content in areas beyond text (e.g., images, video, and code). Table 1 below can help you identify when you need to cite use of commercial, off-the-shelf GenAI outputs and what depth of use is acceptable. Citation of customized solutions may look different. Since the table does not represent the full array of possible use cases, use your professional judgment for your circumstance. If in doubt, cite.

Table 1: GenAI citation guidance

Breadth of Distribution	**Proofreading, Grammar**	**Brainstorming, First Draft, <25% AI**	**Collaborative Writing, About 50% AI**	**Human Edited, >75% AI**	**100% AI Content**
Press release, prepared remarks	✓	Cite	**x**	**x**	**x**
Replies to public inquiries	✓	Cite	**x**	**x**	**x**
Public facing web content	✓	Cite	Cite	**x**	**x**
Memos, broad internal communications	✓	✓	Cite	**x**	**x**
Internal process docs	✓	✓	Cite	**x**	**x**
Source code	✓	✓	Cite	**x**	**x**
Emails	✓	✓	Cite	Cite	**x**

Internal chat messages	✓	✓	Cite	Cite	**x**

✓: No citation needed
Cite: Use the citation template
x: Unacceptable use
*Table adapted from Vermont's Guidelines for Use of Content Generating AI

The standard citation template is as follows; slight variations based on context and type of output (audio, video, image, etc) are acceptable.

"This content was [drafted, edited, translated, created] with the assistance of a generative AI tool [ChatGPT, Gemini, Claude, etc]. The content has been reviewed and verified to be accurate and complete, and represents the intent of [office, department, the State, or person's name]."

V. Prompt engineering

The prompt is the current interface for many commercially off-the-shelf (COTS) GenAI tools. The structure of the prompt goes a long way to improving the results. Ensuring the right type of prompt for your goal is known as "prompt engineering." We suggest reviewing the following resources to help you craft better prompts and improve your outputs. A forthcoming training module will also help you here.

- OpenAI Prompt Engineering Documentation
- InnovateUS module on Prompt Engineering on MS CoPilot
- AWS guide to prompt engineering

VI. Example uses

Below are several example usage scenarios. These are only a small fraction of what's possible, and a forthcoming training module will help you further understand the art of the possible and provide examples. Experimenting with different use cases, using different types of prompts, in different tools, will help you gain an intuition of where these tools are useful and not useful. All guidelines in Section IV above apply to each of these example uses, and in time, we may also ban certain uses, assuming unacceptable risks in light of our AI Principles.

- **Data analysis**: Interact with public datasets through inquiries on tools like OpenAI's Code Interpreter, or generate code to run data analyses.
    - DO:
        - Understand the functioning of code and its logic before its practical application.
        - Leverage GenAI tools as an initial aid.
        - Ensure you are well-acquainted with novel libraries and dependencies. Familiarize yourself with potential vulnerabilities and security aspects related to the chosen language or library.
    - DON'T:

        - Insert PI or sensitive information into prompts.
        - Use code in a production setting without a comprehensive understanding of its operations.
- **Drafting documents**: GenAI tools can help create the first drafts of memos, letters, job descriptions, and other administrative documents.
    - DO:
        - Carefully edit and review any generated content.
        - Leverage disclosures as appropriate (see Table 1).
        - Be as specific in your prompt as you can for better results, and leverage additional prompts after initial generation to further craft the results.
    - DON'T:
        - Use GenAI to create content on sensitive topics.
        - Rely on GenAI to produce accurate information - you must always fact-check.
- **Generate synthetic media**: Leverage prompt-based GenAI tools like DALL-E or Midjourney to create images, audio, or video.
    - DO:
        - Ensure all content aligns with your department's existing web, design, communications, or other relevant standards and policies.
        - Leverage GenAI tools to create mock-ups that facilitate conversation with creative professionals.
    - DON'T:
        - Assume that outputs of genAI are respectful and non-offensive - test and validate with others prior to using.
    - Resource: InnovateUS module on using image generation GenAI tools
- **Improve existing text**: Use GenAI tools to improve existing text in various ways.
    - DO:
        - Leverage GenAI tools to simplify/clarify existing documentation, rewrite in certain styles, specify reading levels, improve grammar, shorten to a desired length, and/or change the tone.
        - Review the updated copy to ensure it is accurate and captures everything from the original text that you need it to.
    - DON'T:
        - Include PI or sensitive information in the prompts.
    - Resource: InnovateUS module on using GenAI to De-jargon government language
- **Summarization**: Prompt GenAI tools to summarize long documents that you upload, and draw out the salient points
    - DO:
        - Be aware that the resulting summary might have biases, as it will tend to present language that is more frequent in the data used to train the model.

        - Particularly if you plan on making a decision based on the summary, you should read the entire document(s) to make sure you do not miss or mischaracterize the original document.
    - DON'T:
        - Upload or include in prompts any PI or sensitive information.
- **Translation**: Use tools like ChatGPT, Gemini, and others to translate any string of text into many different languages.
    - DO:
        - Ask a native speaker to review outputs to ensure accuracy.
        - Use different tools with the same prompt to compare results.
    - DON'T:
        - Include PI or sensitive information in the prompt.

VII. Next steps

DoIT's Governance, Risk, and Compliance (GRC) team is separately building out a standalone approval process to formally review requested tools. We anticipate this process being available later in 2024. Until then, send any inquiries to AI@maryland.gov if you're unsure about the desired tools' use in light of the guidelines in this document.

VIII. Acknowledgements & resources for further reading

The following resources were invaluable in drafting these interim guidelines:

- InnovateUS training modules
- City of San Jose GenAI Guidelines
- NIST AI Risk Management Framework: GenAI Profile (initial public draft)
- State of New Jersey GenAI Guidelines
- State of Vermont GenAI Guidelines
- State of Washington State GenAI Guidelines
- U.S.Department of Homeland Security (DHS) Commercial GenAI Guidelines