Guidance for Artificial Intelligence Tools

There are a great number of uses for technologies that leverage artificial intelligence (AI) and recent developments in "generative AI" — tools that create content like text and images — have generated excitement about new possibilities.

When using any tools, University of Colorado employees are each responsible for understanding how to use the tool effectively, safely, and within existing policies and laws. The goal of the following sections is to educate users of AI tools about key considerations in effective and safe usage.

- How do I get started with AI?
- How do AI tools fit into existing CU policies and procedures?
- What considerations should I have about information put into an AI tool?
- What considerations should I have about the output from an AI tool?
- Examples of situations where generative AI tools might be used at CU
- Generative AI tools available to employees at CU System Administration
- Further reading

Getting started with AI

If you are interested in the possibility of adding an AI tool to your toolbox, walk through these steps to get started:

- Create a list of your goals and requirements
    - Stay focused on the business/functional outcomes and not the tool or path to get there. For example, say "I want my customers to be able to easily find answers to common questions" rather than "I want to use an AI tool to answer customer questions."
    - Mark down which outcomes are absolute requirements and which might be optional
- Document what data or information will need to be used in the process. This will help ensure you have the necessary approval and understand the data sensitivity.
- Review this guidance document for AI considerations to build into your work.
- Be open to both AI-based and non-AI solutions. While emerging technologies provide new opportunities, existing options may still be the best fit.
- Contact UIS for guidance on which options might best fit your needs.
    - Send an email to help@cu.edu asking UIS for AI solution design assistance.

Existing CU policies and procedures

CU has a variety of policies and procedures regarding information technology, information security, data and procurement that may apply to the use of AI tools. CU endeavors to develop policies that apply to a wide range of technologies rather than specific policies about different technologies, and this applies to AI technologies as well.

Universitywide policies

- Information security program policy - https://www.cu.edu/ope/aps/6005
- Data governance
    - Data governance policy - https://www.cu.edu/ope/aps/6010
    - Data governance website - https://www.cu.edu/data-governance
    - Data classification - https://www.cu.edu/data-governance/resources-support/data-classification
- Collection of personal data from students and customers - https://www.cu.edu/ope/aps/7003
- IT procurement processes - https://www.cu.edu/psc/procurement/information-technology-procurement

System-administration specific

- Use of IT Resources policy - https://www.cu.edu/docs/system-administration-information-technology-policy-use-it-resources
- Privacy policy - https://www.cu.edu/privacy-policy

Information put into an AI tool

Typically, the use of AI tools involves a third party handling data provided by their customer. This could be data used to build out the knowledge of the AI system, such as providing it with copies of all your user help documentation so the AI system can be configured to answer questions. Or it could be data that feels more like a question or request, like asking ChatGPT to summarize a long document.

Data sensitivity and approval of use

When using any CU data, it's important to get appropriate approval for the use of the data from the data trustee or steward outlined in CU's data governance process and to understand the sensitivity of the data as described in CU's data classification model. Higher sensitivities of data require stronger protections and might not be appropriate for some types of AI tools.

Contract terms

When using any third-party tools, it's important to understand the terms you agreed to for usage. With AI tools in particular, third parties might include the right to reuse your data to further develop their services. Even if there isn't a formally signed contract, using a generative AI tool likely includes agreeing to some terms and conditions about the data you put into it.

Output from an AI tool

Accuracy of output
Generative AI tools are powerful engines for creating content based on a wide variety of input. The methods used by these tools are highly effective but have limitations and flaws. Whether it's too many fingers on the image of a hand or fictional jobs on a resume, AI tools have demonstrated an occasional tendency to produce inaccurate content. Due to these limits, it's important to have a process for vetting the output of AI tools before relying on it for business decisions or publishing the content publicly.

Safety of generated code
Reviewing the output is especially important when using AI tools to create scripts and programs. Modern tools have demonstrated success in creating code based on a wide variety of content available on the web, but the source content might contain flaws that find their way into the output. These flaws could lead to functional problems or security vulnerabilities. Whether code is generated by AI, written by hand or borrowed from development communities, CU employees are responsible for the effects of code they run on CU systems.

Biases in output
In addition to reviewing the accuracy of content created by generative AI tools, you should also review the output to ensure it meets CU expectations for being thoughtful, supportive and inclusive. Because generative AI tools often build upon content from across the internet, these tools can sometimes reflect biases or even offensive content. For example, Google has made multiple adjustments to its translation tool to remove possible inappropriate output from the system. Biases in output could be more subtle – maybe an image generation tool tends to create people of a particular age, skin tone or gender, or perhaps it highlights stereotypes. These biases can be corrected with thoughtful writing of the requests made to a generative AI tool.

Attribution
You should always be transparent when content is sourced from an AI tool. Include a note or banner on AI chatbots, AI-generated documents and other output from generative AI tools.

Examples of AI-tool uses

Using an AI-driven chatbot to answer customer questions

A popular use of generative AI is to build a tool for answering basic customer questions that are covered by existing documentation. In this scenario, departments might work with their IT staff or a vendor to configure an AI tool to "learn" their documentation and tune the responses given by the tool. A well-developed "chatbot" can be effective at answering common questions and freeing up staff time for more detailed customer needs.

In this example, it's important for a department to consider the three areas listed above:

- Existing policies and processes
    - Follow all procurement processes for an IT purchase, including a security review.
    - Understand the data classification of the information being loaded into the tool.
    - Understand any privacy implications for users of the tool.
- Information put into the tool
    - Understand how the third party handles CU data and what rights they have to use it.
    - Work with security teams to ensure proper controls are in place for the tool.
- Output from the tool
    - Perform lots of testing to ensure the tool provides accurate output.
    - Understand how to adjust the tool if the output is inaccurate.
    - Understand if you have options to review all of the output given to customers and perform a periodic review.

Using a generative AI tool to summarize text or data

Artificial intelligence tools have proven useful in summarizing information and are commonly used to write summaries in many fields. This ranges from writing a two-sentence summary of a news story to generating a list of the top ten polka songs from a wide variety of input. Let's visit our three areas again with this example:

- Existing policies and processes
    - Follow all procurement processes for an IT purchase, including a security review.
    - Understand the data classification of the information being loaded into the tool.
    - Understand any privacy implications for users of the tool.
- Data sensitivity and handling
    - Understand how the third party handles CU data and what rights they have to use it.
    - Work with security teams to ensure proper controls are in place for the tool.
- Vetting the output
    - Any summaries should be reviewed by someone who is strongly familiar with the original content, to ensure the key points are captured.
    - Be sure to include an attribution that the summary was created by an AI tool.

Using a generative AI tool to translate information into another language

Translating content into additional languages is another popular use for AI tools. This can open access to content and services to more individuals and be more inclusive of community members whose native language differs from the language of your content.

- Data sensitivity and handling
    - Ensure that the information to be translated fits within the approved data classification for the tool.

- Vetting the output
    - Be sure to have the output reviewed by an individual who is fluent in the output language, preferably a native speaker of that language. Native speakers can catch a culturally inappropriate use of words that might be missed by others.

Further reading

Artificial Intelligence Risk Management Framework from the National Institute of Standards and Technology (NIST AI 100-1) - https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf
Microsoft site on responsible AI - https://www.microsoft.com/en-us/ai/responsible-ai

Google site on responsible AI - https://ai.google/responsibility