IWU Generative AI Policy

This document provides guidelines for the use of Generative Artificial Intelligence tools ("GenAI" in the context of this document) by employees of IWU, including both faculty and staff. GenAI includes text generation tools (large language models such as OpenAI's ChatGPT, Google's Gemini, and Anthropic's Claude), media generation tools (including image, audio, music, and video generation), AI agents and agentic systems, and systems that incorporate those tools into other applications. This document outlines the risks involved in using this technology, places limitations on when and how they should be used, and offers guidance on how to use them appropriately.

Guiding Principles

1. Above all, the goal of these policies is to avoid harm to individual members of the institution as well as to the institution itself.
2. GenAI encompasses novel technologies that are developing quickly. This document attempts to be forward-looking, but some of its policies must remain non-specific to accommodate advances in the state of the art.
3. As GenAI is an information technology, IWU's existing ITS policies and procedures apply to all uses of GenAI.
4. IWU and its employees must remain in compliance with relevant laws and regulations such as the Gramm-Leach-Bliley Act (GLBA) and the Family Educational Rights and Privacy Act (FERPA).

The Technology

GenAI systems are large neural networks (or other "machine learning" systems) often called "models" that are "trained" to generate realistic data in some form (text, image, audio, music, video, etc.). They are trained using huge amounts of existing data of whatever form they are intended to output, such as text from books, journals, newspapers, and the world wide web, in the case of text generation. Once a model has been trained, it can be used to generate new data of that form based on some prompt from a user, often in the form of a written question, description, or uploaded file. The generation process is a complex mathematical computation using all of the information processed and compressed into the model during the training process.

AI agents and agentic systems go further, applying GenAI technology to drive real-world actions (such as interacting with websites, managing files, sending emails, or running code) rather than simply presenting generated content to the user.

Costs and Risks

All uses of GenAI should be undertaken with the following in mind.

Privacy and Ownership

Nearly all GenAI tools are developed and hosted by corporations external to IWU. Using such a tool requires sending one's prompt to that organization, at which point they have full access to and control over that data. Prompts often include information one wants to incorporate into the generated result in some way. Each company has its own policy regarding how it might use and retain that information, and there is little to no external oversight of that. Some companies explicitly state that they may use any information submitted to them when training newer models.

Confabulation and Falsehood

While GenAI models are trained to output realistic, accurate content (whether text or otherwise), they are statistical processes with no guarantee of correctness. Models are well known for "confabulating" or "hallucinating," sometimes presenting false information as true with confidence and no indication of its falsehood. The frequency and scale of confabulation changes based on the amount of information in the training data relevant to what is being generated. It is difficult to predict how likely any generation is to be correct and accurate. This becomes substantially more dangerous in AI agents, as their actions are unpredictable and may be misaligned with the user's intent and interests.

Embedded Biases

GenAI models are trained on large corpora of existing media, and they incorporate the biases embedded in that media. Companies training models take measures to remove and reduce those biases, but some always remain. These biases can include biases against protected classes, such as gender, sexual orientation, race, and religion. Those biases can affect and be evident in the generated output of a model.

Environmental Impact

The computation required to train a single GenAI model is substantial, using large amounts of electricity and water (for cooling) continuously over periods of months, each costing potentially tens or hundreds of millions of dollars over that period. Applying a model to generate a single output takes much less time, but each new application requires an additional amount of energy and cooling.

Copyright and Intellectual Property

The data used to train GenAI models typically contains huge amounts of copyrighted material generated by people who were not aware their work would be used in this way. The legal rights to use copyrighted material in this way are still unclear (under litigation beginning in 2024), and the moral rights to do so are strongly disputed by many.

Perceptions of Generated Materials

For many of the reasons above as well as others, many people have a strong negative perception of materials generated by GenAI and, by extension, those who use or publish such material. Always consider the audience and whether they may perceive work created using GenAI as being tainted, low quality, and/or immoral.

Policies

1. Employees (faculty and staff) should consider all of the above costs and risks when deciding whether using GenAI is appropriate in any new context.
2. Employees must not include sensitive, confidential, or personally identifying information in any prompt or other input sent to an external or cloud-based GenAI tool. This includes but is not limited to financial information, health information, and academic records (including grades).
3. Employees must not connect AI agents or agentic systems that have not been previously approved by ITS to any IWU systems, data, or accounts via sharing credentials or any other mechanism.
4. All output of any GenAI tool must be reviewed by a human for accuracy, bias, and suitability before using, sharing, or publishing that output.
5. Employees should consult with their departments and/or professional organizations regarding standards and conventions for the use of GenAI in their field.
6. Within the constraints of the above, individual faculty have discretion over how they and their students may use GenAI within their classes.