An Initial Guide to Generative AI at WIU:
The Fall 2023 AI Task Force's Final Report to Faculty Senate

Committee members:

Michael Lorenzen (Chair, Libraries)

Everett Hamner (Lead Writer, CAS)

Graciela Andrango (CBT)

Daniel Atherton (CBT)

Rashmi Sharma (COEHS)

Gregory Kain (ex officio, Executive Director of University Technology)

Presented to:

Faculty Senate

Purpose, process, and foundational principles

The Faculty Senate AI Task Force was charged with examining how WIU faculty can best integrate AI technologies into their teaching (where suitable) and best prepare students to use AI appropriately (rather than unethically).

During near-weekly discussions during the Fall 2023 semester, we quickly recognized a shared understanding that as AI capabilities continue to advance, it is critical for all of us to be proactive about both the opportunities and dangers presented. While this report was occasioned by recent expansions of generative AI (genAI) and Large Language Models

4) *Since employers w*

Definitions, subcategories, and timelines

Part of why we needed an ad hoc committee to build this report is that there are m

natural language processors (NLPs), which enable verbal human-computer interaction, including the algorithms behind many chatbots and language translators;

increasingly accurate and localized weather predictors;

various tools of facial recognition, image analysis, and photo tagging;

ad targeters and streaming platforms (like Netflix) that create tailored recommendations based on users' historical data; and

virtual assistants like Siri or Alexa.

Some thinkers classify *generative AI (GAI* or *genAI)* as a variety of ANI, while others argue that it represents an intermediate step between ANI and AGI. Many of the tools just listed as ANI are being transformed further by genAI, and the lines get blurry. It's probably simplest to say that presently, the vast majority of researchers agree that publicly-released genAI represents major steps forward – some emergent elements of which we are still struggling to understand – but also resist suggestions that it shows significant signs of consciousness or self-awareness, which are required in most definitions of AGI.

What makes AI generative? A simple rule of thumb is that its primary purpose is innovation and creativity (even if some would qualify these terms' application). All genAI relies on *machine learning*, a subcategory of AI that dates back many decades and that involves *optimizing* processes. Machine learning has gradually taken AI beyond classical symbolic approaches repeatedly following pre-defined protocols (like *nearest neighbors* and *decision trees* models) and into connectionist approaches involving more obscure processes. Sometimes contemporary genAI requires enormous datasets featuring laborious labeling of included items, but increasingly, new forms can intake data in relatively "raw" forms.

Artificial General Intelligence (AGI). The broken half-line above signifies, "we're not here yet – not in any public form, anyway." But we're very likely headed here, so it's important to understand that AGI refers to AI that can achieve a very wide range of human tasks at or above average human ability levels, perhaps via robots that can agilely navigate unpredictable physical environments. Most definitions of AGI assume consciousness, the criteria for which inspires its own debates among cognitive scientists, philosophers, and others. Some of us will want to stress ways in which present forms of AI fall short of AGI, and that is important to understand. At the same time, we should observe that the boundaries between human and AI capabilities keep falling, and at what seems an accelerating pace. In March 2023, GPT-4 earned a 90% on the bar exam, aced the GRE (99% on verbal, 80% on quantitative), and earned a 4 or 5 on nearly every AP test. In recognizing how these goalposts are culturally and historically defined, we need to keep asking: who gets to decide what constitutes human-level intelligence, and on what bases can they make this determ]

Opportunities and dangers

Your

On-demand manufacturing processes, personalized robotic assistants, and efficient supply chains

Individualized market analysis, portfolio management, and stock trading services

Cybersecurity, e-commerce, mobile data processing, and fraud detection

These can be very good things; in many contexts, they already are. Of course the devil is in the details, but many people are willing to allow some measure of information sharing or surveillance in exchange for what can be enormously liberatory prosthetics for those living with disabilities, game-changing small-business solutions for those with minimal start-up capital, and situationally specific legal and medical advice that would normally be too expensive for people living in less developed nations.

Conversely, for our colleagues who are more immediately enthusiastic about genAI, we want to drive home the potential for it to be abused and to yield unintended or unadvertised consequences. The same tools that can serve life and justice can be means of exploitation, consciously or unconsciously. Here are brief looks at some of the most recognized weaknesses and dangers in some current forms of genAI:

Algorithmic bias. Remember, generative AI is predicting desired output on the basis of affirmed input, i.e., its *training sets*. It relies on enormous datasets to establish its (incredibly) educated guesses, but whose predilections and expectations are most represented in those datasets? In too many cases, inputs overrepresent the perspectives of longstanding majority groups (those who are White, male, wealthy, straight, cis, able-bodied, and Christian, just to name the most influential groups in a partial list of identity categories). The timeworn adages of athletic trainers and nutritionists everywhere applies reasonably well to AI: "garbage in, garbage out," and "you are what you eat." That is, poorly tuned and/or unrepresentative data inputs leads to generative AI outputs that inevitably perpetuates ugly patterns.

Here's a simple, oft-cited example in the journal *Science* from 2019. As explained in Ziad Obermeyer et al.'s abstract, "Health systems rely on commercial prediction algorithms to identify and help patients with complex health needs. We show that a widely used algorithm, typical of this industry-wide approach and affecting millions of patients, exhibits significant racial bias: At a given risk score, Black patients are considerably sicker than White patients, as evidenced by signs of uncontrolled illnesses. Remedying this disparity would increase the percentage of Black patients receiving additional help from 17.7 to 46.5%. The bias arises because the algorithm predicts health care costs rather than illness, but unequal access to care means that we spend less money

Susceptibility to error and hallucination (or "hallucitation"). Because genAI is all about predicting the next word or phrase that seems most likely or appropriate, and because it in no way "understands" the output it is producing, it too often sacrifices accuracy for speed. Lacking a scholarly reference, it may cite a researcher as having claimed something they never wrote, sometimes in an article that does not exist. It might make up a fictional scholar or journal, if that most closely approximates what its weighting system grades the next most appropriate word or phrase. Sometimes the results are humorous, but they can also be very dangerous – especially in a time when many citizens are already struggling to recognize political disinformation.

Lack of transparency. Ironically, the more that genAI is utilizing multi-layered neurological processes that in some ways emulate our own, the less we are capable of understanding its means for producing outputs. This is the "black box" problem: as companies feed enormous amounts of training data (remember, terabytes, petabytes…) into AI systems and increasingly rely on "unsupervised" models, it is becoming impossible for human observers to witness or grasp the processes by which system optimization and heightened prediction accuracy is achieved. This might mean greater ease in gaming these systems without detection; it also becomes easier for some companies to rationalize enormous "mistakes" as the cost of progress. It is also worth noting in this

[Phone Number]
[Date]

MIT Press
Attn: Book Requests
55 Hayward Street
Cambridge, MA 02142

Subject: Request for Exam Copy of "Robot-Proof" for English 347 Course

Dear MIT Press Book Requests Team,

I hope this letter finds you well. My name is [Your Name], and I am an instructor at [Your Institution], currently teaching an English course titled "The Novel in Context: Artificial and Human Intelligences" (English 347). I am writing to request an exam copy of the book "Robot-Proof" for consideration as a text for my upcoming course in the spring semester.

Book Details:
Title: Robot-Proof
Author: Joseph E. Aoun
ISBN: [Please look up the ISBN online for the most accurate information]

After researching potential texts for my course, I have identified

How about a more complex task? What if one of us were to use genAI to draft *this section* of our report about teaching applications? After submitting bullet points from our committee's conversations, we quickly received and lightly edited the following genAI output:

The integration of AI [6]

AI can be a transformative tool in modern education, augmenting human capabilities and enriching the learning environment.

Of course neither this short essay nor the business letter above constitutes perfect writing. Both are grammatically clean and reasonably organized, but many of us would revise them differently. For example, they exemplify genAI and especially LLM tendencies to feature unnecessary repetition. After all, LLMs are prediction engines, using seemingly innumerable training documents to reasonably estimate the most likely "next word" in a given sentence, given what has come before. These tools don't "understand" what they are saying, so they are prone to say the same thing repeatedly in multiple ways, especially when user prompts lack word limits or do not explicitly prioritize brevity.

Nonetheless, for many, such tools offer a helpful way of getting past the intimidation of the blank page. They will not approach many writers' *ceilings*, but LLMs can get some writers off the *floor*. Of course many students will be sorely tempted to simply provide a prompt, quickly touch up genAI output, and submit the result, and where this unethical usage pattern is detectable, we expect it will continue to require active monitoring and correction. But we hope the examples above suggest how an AI's initial composition or proposed outline could also provide a helpful starting point or template for more thorough revisions, and how after a writer has generated their own draft, they might benefit from asking genAI to generate a comparable draft or otherwise direct it

Syl

of the requirements of students in *genAI permitted courses* also apply in these settings. Some examples of specifically allowed or disallowed uses might include:

> brainstorming and refining ideas;
>
> strengthening research questions;
>
> seeking additional information on a research topic;
>
> drafting an outline from notes or to rethink one's initial organization;
>
> composing an initial draft or an alternative draft with which to compare one's own;
>
> writing full, unedited sentences, paragraphs, or papers; and
>
> checking grammar, style, and other mechanical issues.

Faculty may also identify particular course contexts in which genAI may or may not be used, e.g.:

> online course discussion boards;
>
> group projects;
>
> quizzes or tests, whether in-person or take-home/online; or
>
> specific papers.

A *genAI partially permitted course* might include a syllabus statement like the following:

> This is a *genAI partially permitted* course.
>
> For Papers 1-3, the use of genAI tools like ChatGPT and DALL-E for course deliverables are welcome. However, please note that this must be explicitly disclosed after the conclusion of each assignment usage as follows: "This paper/take-home exam was completed with the assistance of [genAI tool] in the following ways: [outlining/drafting/mechanical polishing/etc.]." In making all such submissions, students affirm that they have used genAI to expand rather than replace learning; that they have fact-checked its compositions; that they have protected the privacy and intellectual rights of all members of our learning community; that they accept full responsibility for all elements of their submission; and that they are prepared to provide complete records of genAI usage via screenshots showing date/time.
>
> For all discussion board postings, quizzes, and tests, genAI is not permitted. In these cases, students will interact with each other and faculty in order to develop skills independently of AI tools and to demonstrate independently achieved insights and express attitudes in a more personal voice than genAI might convey.

If any other potential uses or gray areas arise, the burden is on students to consult faculty for explicit clearance to utilize genAI tools. If in doubt, ask!

do not allow this technology's usage in any course assignments and tasks, unless specifically and explicitly identified for a very limited context such as a particular

pass along as a potential model for faculty wishing to be especially clear about their expectations on a given assignment:

By submitting an assignment for evaluation:

you assert that it accurately reflects the facts and to do so you need to have verified the facts, especially if they originate from generative AI resources;

you assert that all your sources that go beyond common knowledge are suitably attributed. Common knowledge is what a knowledgeable reader can assess without requiring confirmation from a separate source;

you assert that you have respected all specific requirements of your assigned work, in particular requirements for transparency and documentation of proces

Preparing for the future of AI at WIU and beyond

This report's most central points about the present apply equally to the future: we must keep assessing particular AI applications' strengths and weaknesses, their pr

Cantor, Matthew. "Incredibly smart or incredibly stupid? What we learned from using ChatGPT for a year." *The Guardian* (October 12, 2023).
https://www.theguardian.com/technology/2023/oct/12/chatgpt-uses-writing-recipes-one-year

Chan, Cecilia Ka Yuk. "A comprehensive AI policy education framewu,

Heaven, Will Douglas. "Welcome to the new surreal. How AI-generated video is changing film." *MIT Technological Review* (June 1, 2023).
https://www.technologyreview.com/2023/06/01/1073858/surreal-ai-generative-video-changing-film/

Heikkilä, Melissa. "Making an image with generative AI uses as much energy as charging your phone." *MIT Technological Review* (December 1, 2023).
https://www.technologyreview.com/2023/12/01/1084189/making-an-image-with-generative-ai-uses-as-much-energy-as-charging-your-phone

Hutson, Matthew. "Will superintelligent AI sneak up on us? New study offers reassurance." *Nature*