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

Western Illinois University

29 December 2023

Contents

Purpose, process, and foundational principles	3
This report-guide aims to help WIU faculty proactively address both opportunities and challenges associated with AI, especially Generative AI (GAI or genAI) and Large Language Models (LLMs).	
Definitions, subcategories, and timelines	5
We start by distinguishing genAI and LLMs from, say, the robots of Terminator 2.	
Opportunities and dangers	8
Here we survey some of genAI's most appealing applications and detail some of its dangers (e.g. via algorithmic bias, lack of transparency, job loss, susceptibility to error, or nefarious applications).	L
Integrating genAI applications as teachers	12
Utilizing two examples of genAI output, this section considers specific genAI strengths and weaknesses for particular kinds of faculty work life and classroom teaching contexts.	
Syllabus policy options and recommendations for helping students use genAI responsibly	16
We recommend faculty develop unique syllabus statements within three broad genAI usage/non-usage categories and offer our recommendations to promote ethical genAI usage at WIU.	ge
Preparing for the future of AI at WIU and beyond	20
In this final, most speculative and interrogative section, we look ahead at some of the key questions we expect to shape the development of genAI and other AI technologies and their impacts at WIU.	

Appendix: Extended bibliography

We offer a recent-article-focused bibliography subdivided into pedagogical discussions, larger higher ed and theoretical issues, and even broader ethical and social questions concerning AI's evolution.

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 (LLMs), especially in such popular forms as DALL-E 2 and ChatGPT-3.5 and 4, our conversations took up these developments with an eye to what is now over half a century of progress around various forms of artificial intelligence and machine learning.

This report considers a wide array of areas that must be considered if WIU is to foster informed and ethical usage of generative AI especially. If our conversations this semester taught us anything, it was that we would not be resolving all of the philosophical or pedagogical questions these technologies raise or freshly necessitate. While we regularly benefited from feedback via a range of WIU faculty and staff not serving on the Task Force, we also expect that there are more perspectives and experiences among our colleagues that are not yet represented here.

Several foundational principles shape our recommendations across this report:

- 1) We expect that AI capabilities will continue to change rapidly. Anyone familiar with previous rapid expansions in a particular technology will understand the dangers involved in excessively confident or specific predictions. We all understand how cell phones today make toys of those from the early 2000s and how the personal computers and laptops of the 1990s are bulky and slow by today's standards (even as they seemed unimaginably compact and efficient to the scientists who used room-sized devices to put the first human beings in orbit). We expect similar advances in AI over the coming months and years. In fact, some scholars expect history to compare AI's expansion in this period to the development of the first atomic bomb (though hopefully with less violent connotations). While offering a historical review and some very loose projections below, we would advise caution about ruling out particular AI capabilities. With a similar eye to maximizing the report's shelf life, we avoid focusing too fully on any one commercial AI application.
- 2) While the release of ChatGPT in November 2022 particularly brought AI to many faculty's attention, genAI is about much more than composing essays and other written documents. GenAI can also create visual images, videos, music, spreadsheets, and computer programs. It can design web pages, create citation tools for scholarly articles, answer email, and maintain personal calendars. There are risks involved in such applications, but the point is that AI can automate a great deal of work that can be done on a computer, at widely varying levels of quality and cost.
- 3) The Task Force does not regard AI as inherently good or inherently evil. Our shared approach aims to be (a) moderate, recognizing that most technologies bring both new opportunities and dangers; (b) pragmatist, accepting that AI will likely continue to heavily impact human cultures for the foreseeable future; and (c) critical, emphasizing the opportunity this technology presents to clarify the processes of critical thinking and the pursuit of reliable knowledge that are at the heart of our shared academic mission. Just as the cultures we engage are dynamic, so we all need to continue adjusting our teaching and learning; at the same time, there are major concerns to keep in view.

- 4) Since employers will increasingly expect WIU graduates to use generative AI in the workforce, we need to teach students to utilize these tools thoughtfully. In some contexts, this learning will involve practicing new skills by working directly with various instances of generative AI, which are quickly being integrated into many common software packages. In other cases, this learning will exclude direct AI utilization, but may still involve thinking about the strengths and weaknesses of automation, algorithms, and machine learning, especially as they compare to human abilities and cognitive processes. We also expect there to be many hybrid approaches.
- 5) Taking this period of transformation seriously means continuing to foster students' abilities to think, write, create art, code, and otherwise shape the world independently of AI. Being able to use any tool critically and creatively requires a deep understanding of its unique capabilities and drawbacks, and in some contexts, that is only achievable via extensive practice of a skill in the *absence* of said technology. We expect that some programs and disciplines will prioritize direct generative AI adoption while others will emphasize skill development in isolation from these tools, and we believe students should experience a variety of approaches.
- 6) One of the most common and immediate concerns expressed by faculty over the last year has been about unethical applications of generative AI in writing. We address this much more extensively below, but our primary recommendations for faculty are these: (a) to understand the capacities of the technologies in question as fully as possible; (b) to resist the temptation to assume that any so-called "AI detector" software can be relied upon to identify writing produced with AI assistance (note especially the number of false positives, the disproportionate flagging of non-native English speakers, and the destruction of trust that comes with misguided accusations); (c) to prioritize clear expectations for student AI use and/or non-use, and (d) to consider emphasizing process development even more fully, including in grading rubrics.
- 7) Ultimately, the Task Force wishes to strongly affirm that *the best ways to approach AI will be different for each discipline, course, and instructor.*

Definitions, subcategories, and timelines

Part of why we needed an ad hoc committee to build this report is that there are many overlapping terms involved and their familiarity and usage differs massively according to discipline. For the purposes of this report and for basic understanding of our subject matter, though, let's start with these:

- Intelligence. As many readers of this document could explain in great detail, this term has often been used, especially during and since the Enlightenment, to illuminate diverse habits of logic and reason in building knowledge. At times these processes have been defined against emotion and subjectivity, but there is growing understanding that intelligences are multiple, culturally shaped, distinctly embodied phenomena.
- Human intelligence. Psychology, cognitive science, and other disciplines increasingly underline the need to understand human neurological processes in relationship with those of a wider range of animal and plant species, not to mention digital entities. In a time when our colleagues are demonstrating the capacities of other primates to recognize individuals after decades-long separations – or of trees to use root networks to cooperate against disease threats over miles-long distances – it is well worth reconsidering human intelligences alongside those of other species we may have previously overlooked.
- AI (Artificial Intelligence). Once we focus on artificial intelligences, we are generally describing silicon and circuit boards rather than organic material and brain layers. AI's development has taken place over many decades, not just the last couple years. The earliest instance of AI is debatable, with some historians pointing back centuries and focusing on conceptualization, not just physical instantiation. However, most historians of science point to the foundational significance of the mid-1950s coining of the term "artificial intelligence" and programs and networks like 1955's Logic Theorist and 1957's Perceptron. These were the first tangible examples of AI to begin fulfilling science fictional and fantastic visions of artificial intelligence like those in the play *R.U.R.* (1921) (which introduced the term "robot"), the films *Metropolis* (1927) and *The Wizard of Oz* (1939), and short stories by Isaac Asimov (1940ff.).
- Artificial Narrow Intelligence (ANI). Sometimes also called "weak AI," this term is generally juxtaposed against Artificial General Intelligence (AGI) or "strong AI" (which is described below). Though there are always debates about boundaries, ANI usually refers to our present stage of publicly known artificial intelligence, whereby the tools in question are capable of specific, limited skills and command fulfillments rather than a wider range of actions and multifunctional tasks akin to those routinely doable by most human beings. (Note the potential for ableist assumptions to influence such discussions!)

Current ANI applications include:

- chess and Go players (which have consistently demonstrated their capacity to beat humans, although human and AI teammates can still beat AI operating alone);
- self-driving vehicles (ranging from adaptive cruise control to fully autonomous, steering-wheel-less cars);

- 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" form, organizing inputs itself rather than necessarily relying on human supervision. This is possible via machine learning innovations in *neural networks* and *deep learning* that emulate the multi-layered processes of human brains (even if they sometimes generate strategies that seem unfamiliar and counterintuitive to our mental processes).
 - Large Language Models (LLMs) are a subcategory of genAI focused on the production of language (rather than, say, images). While there is no precise determination of what constitutes "large," they utilize written language inputs at quantities a million times larger than those used a decade ago. For non-computer scientists, all we need to know is that with the most recent examples, however many gigabytes of input we're imagining is likely too small. This is why terms like "terabyte" and "petabyte" exist....
 - Another, partially overlapping term is *foundation models*. These are the underlying models that can be modified for new purposes. E.g., OpenAI used GPT (Generative Pre-Training Transformer) 3.5 as the foundation for the chatbot that became ChatGPT.

- 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 determination?
- Artificial Super Intelligence (ASI). Again, the visual barrier extends only halfway across the page, but this time it's a solid line. There are more philosophical and technical debates about the possibility of ASI than AGI; now we're talking about "the Singularity," a point at which futurists imagine AI becoming not just equivalent to human capacities (AGI), but godlike in its ability to control society. At this stage, most thinkers assume AI is not just self-aware but so powerful as to be capable of saving or destroying humanity. While it's unwise to fully rule out the kinds of scenarios dramatized in films like Terminator 2 (1991), I, Robot (2004), or Ex Machina (2014), faculty should understand that the genAI discussed in this report remains quite distant from such possibilities.

Opportunities and dangers

Your conference is in a new city this year, one without an easy train straight to downtown. You get off the plane (finally) and make your way to the terminal exit, and it's late. There's a bus, but it won't come again for 45 minutes.

If the year is 1994, there's a long line of cabs waiting. State your hotel name, and that's it. The cabbie knows the place, plus five different routes. (But they might take the one creating the largest fare.).

If it's 2024, you could still take a cab, but there aren't so many, and they're even more expensive. Instead, you hop on your phone and in fifteen seconds have competing Uber and Lyft bids, with multiple vehicle sizes and comfort levels. You can prioritize an EV and limit yourself to cars already within a 2-minute radius. Momentarily, you are seeing a picture of your driver, an image of their vehicle, its license plate number, and its real-time location. You pay and tip through the app too, with the funds automatically transferred to the driver's account, minus a slice for your virtual matchmaker.

Is such use of AI good or bad? Our answer: both, and it depends on your priorities.

On one hand, Lyft and Uber are super-convenient, especially once you're accustomed to their apps. These services also feel relatively safe: there is a live record of both your and your driver's movements, and since you review each other, those who abuse the system quickly develop poor reputations. The system also has efficiency advantages over taxis: drivers guess less about when and where riders may appear, so there is less idle time, and because nearby vehicles get priority, there are fewer empty vehicle miles.

On the other hand, we lose something human as life-long cabbies disappear. They know their cities uniquely, sensing when something is out of place in ways digital-map-reliant part-time drivers cannot. Some urbanites also miss cabbies' expertise; when you ask, "is that restaurant good?", you don't get a travel site's average rating, but personal advice. Yet potential remains for abusing the immense data involved in AI-dependent systems like Uber and Lyft: who receives surveillance access to their data? Whose insurance rates are affected? Do these systems prioritize wealthy customers and their larger tips over service to others?

Our priority in this document is genAI in higher ed, but we provide this simple example of another common AI application today to suggest how our advice is inevitably shaped to fit particular values at specific moments. It's also worth noting how our cumulative choices shape whole systems. But now let's add to the picture by surveying a wider range of genAI's strengths and weaknesses, blessings and curses, and opportunities and dangers.

For those most hesitant about generative AI, we want to highlight its already widespread benefits. Like AI more broadly, this technology excels at broad and rapid searching and indexing and at pattern recognition, but now it is capable of rapidly generating new compositions of words, images, and code. It can not only translate information across languages and contexts, but subtly modify its output to fit a user's preferred tone, style, and length expectations. In addition to the current applications listed on page 5, the use cases are rapidly expanding. The same tools that do automatic spelling and grammar checks and that offer us sentence-completion suggestions in email and messaging apps are now expanding global access to:

- > Medical imaging analysis, real-time surgery enhancement, and innovative pharmaceuticals
- > Legal discovery tools and case law research in a wide array of languages
- > Live location searching, routing, and complex system scheduling

- > 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 caring for Black patients than for White patients. Thus, despite health care cost appearing to be an effective proxy for health by some measures of predictive accuracy, large racial biases arise. We suggest that the choice of convenient, seemingly effective proxies for ground truth can be an important source of algorithmic bias in many contexts."

This example is telling: the problem isn't that the data is *wrong*, but that it is *misapplied*. An assumption is made that proves badly misguided: while it might be easy for a relatively privileged researcher to assume that less money being spent on health care means less need exists, that is hardly the case. As a result, what seems an unquestionable high-tech presentation of "cold hard data" via a sophisticated AI application only exacerbates long-standing patterns of racial disparity in U.S. health care access.

- 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 questions around copyright, privacy, and consent that genAI raises; the 2023 screenwriters' and actors' strikes highlighted just one very visible context in which the rights of artists and others are vulnerable to irresponsible, unregulated AI training.
- Disproportionate job-loss impacts. Insofar as racial minorities, lower-class Whites, and immigrant non-citizens make up disproportionately high proportions of the less educated workforce, a likely impact of expanding generative AI usage is wider experience of layoffs and unemployment for less-privileged communities. As Ted Chiang compellingly argues in one of the most consequential articles of 2023, keeping A.I. from becoming "the new McKinsey" must be a goal for anyone who values social justice. At the same time, we should realize that generative AI may eliminate many middle-class jobs, particularly those of employees whose work is relatively formulaic and repetitive. (For more on these projections, see especially Rakesh Kochhar's article for the Pew Research Center in the bibliography.)
- Other potential for nefarious applications. Just as genAI is uniquely suited to crime detection, it can be marshaled for large-scale, incredibly rapid theft and misuse of information. For instance, FraudGPT and WormGPT (creative naming, eh?) are being used to compose targeted malware code and phishing emails that generate alarmingly high rates of clicks and downloads.
- Other social and environmental injustice concerns. The amounts of electricity and water consumed by genAI model training and queries are being researched in earnest, and current and potential future impacts are unclear; there is even more silence around the low-wage labeling work required in training many genAI models. However, early studies are suggesting that we could be talking about the annual energy consumption equivalents of small countries, and the *Wall Street Journal* places genAI energy usage above that of all EVs worldwide (see articles by Foy, Heikkilä, Singh, Mims, and Wells). Meanwhile, there are disturbing pieces of investigative journalism about the willingness of genAI companies to exploit workers who lack attractive alternatives (see the pieces by Dzieza and Perrigo).

For readers of this report with more technical interests, here are further details about distinctions among, capacities of, and vulnerabilities of machine learning algorithms.

Machine learning (ML) algorithms consist of *supervised, unsupervised, semi-supervised* (hybrid of supervised and unsupervised), and *reinforcement* methods (Sarker 2021, 2). ML algorithms make classification, regression, dimensionality reduction and visualization, and data clustering of extremely large datasets possible. Supervised algorithms require the use of training data to facilitate classification of data sets into distinct classes. In addition to classification, supervised algorithms can also be used for *regression analysis* on a data set to predict a result (dependent variable) based on one or more predictor(s) or independent variable(s).

Like supervised algorithms, unsupervised algorithms have the ability to cluster or classify objects into distinct groups. Unsupervised algorithms make use of trends or patterns found in the datasets themselves and do not require a *training dataset*, but in some cases, require a user-determined *cluster constraint*. Unsupervised algorithms can be used in a broad range of areas like cybersecurity, health and behavioral analytics, and e-commerce and mobile data processing (Sarker 2021, 9).

The main difference between most supervised and unsupervised algorithms is the use of a training dataset versus cluster constraints. The utility of the supervised algorithm is dependent upon the correctness of the training dataset, meaning an accurate training dataset is essential for supervised algorithms, whereas a poor training dataset may reduce algorithm usefulness. In addition to training dataset validity, *overfitting* the algorithm to the training dataset can reduce model usefulness. Likewise, for unsupervised algorithms, cluster constraints can affect the usefulness of the algorithms, but newer algorithms have been created to negate this issue.

Unsupervised algorithms also include the capability to reduce dataset dimensionality. Reduction of very high dimensional datasets decreases *model data redundancy*, which reduces model overfitting. One popular technique for reducing high dimensional datasets is *principal components analysis (PCA)*. PCA extracts principal components from the original dataset which represent the variation within the dataset in decreasing order, with PC₁ representing the largest variation, PC₂ representing the second largest variation, and so on.

Reinforcement learning is another common machine learning methodology that can be used for analysis and control of robotics, aircraft, and manufacturing systems as well as logistics and system optimization, image processing and classification, and natural language processing (Sarker 2021, 13-14). A popular ML reinforcement learning technique is *artificial neural networks (ANN)*. There are many ANN algorithms that may be used depending upon the problem to be solved, but each commonly makes use of three main layers; an input layer, an output layer, and a middle layer which is hidden from the user. The hidden layer is made up of one or more layers depending upon the algorithm's methodology. One of the issues to consider with ANN algorithms is a lack of transparency. The hidden layer is essentially a *black box* that does not provide transparency as to how a certain outcome was produced. Reinforcement learning algorithms can also make use of hybrid unsupervised algorithms for reduced dataset complexity like PCA.

All the algorithms mentioned here may suffer from bias or other errors due to poor, erroneous, biased, or insufficient data since machine learning algorithms rely on vast amounts of valid and unbiased data. These algorithms can also be used for nefarious purposes, as noted above.. Improper usage of the algorithms themselves from issues like overfitting and poor training dataset selection can also degrade the usefulness of ML.

Integrating genAI applications as teachers

Why not just get rid of the teacher and student entirely? Just let computers talk to each other all day.

This is not teaching and it's not learning: there's nothing real behind it.

You can't automate caring.

These paraphrases of WIU faculty reflecting on genAI may or may not be familiar, but they express real frustrations. The concerns are particularly salient for those of us whose disciplines treat writing and art as means to *creating* and *discovering* knowledge, not just reporting on it. We are now seeing student paper and project submissions that show every sign of being hastily and thoughtlessly produced with genAI.

On one hand, we must be very careful about assuming we can consistently identify every use of these algorithms, because "false positives" are widespread too, and potentially very damaging to faculty-student trust. On the other hand, when expert wordsmiths receive unsophisticated "copy and paste" submissions from ChatGPT 3.5, the problem can be glaring. How do we adjust assignments and assessment processes so that learning is not short-circuited by such tools?

Another concern: it is easy for appropriate enthusiasm about genAI tutorial capacities to bleed into new forms of overreliance. It is one thing to develop greater familiarity with and interest in a world language, for instance, via Duolingo; it is another to build a deep understanding of how an unfamiliar culture has constructed reality through the complex development of words and grammar across time. For some disciplines, writing and creativity are secondary to number-driven analysis, so quite naturally words and images function more technically and transactionally. Fields that revolve around writing and art, however, pay much greater costs when genAI is used to *replace* rather than *supplement* creative processes (not to mention the experts committed to them).

Still, such overreaches should not obscure the fact that genAI *can* be a valuable assistant to many of us. Just because there are real limitations in what any automated system can achieve doesn't mean its gains are illusory. Even in writing-centric fields, for instance, it may help some students and faculty to use AI to generate a "template response" for subsequent revision, or perhaps to utilize this output to reveal areas for potential expansion in a draft they have already composed.

Meanwhile, perhaps especially for faculty and students in number-centric fields, it's worth noting the potential time savings for relatively routine tasks. Here's a simple example. As professors, we must stay current on textbooks in our fields, so when we run across an intriguing new possibility, we request that the publisher send an exam copy. Some distributors use fill-in-the-blank online forms, but others request a letter on university letterhead. One easy option is just to pull up the last such letter we wrote, modifying as needed, but what if we don't have that handy?

It isn't so different to prompt an LLM as follows: "Draft a one-page business letter to MIT Press requesting an exam copy of the book *Robot-Proof*. Explain it's for an English 347 course next spring, The Novel in Context: Artificial and Human Intelligences." Here is an instant response to that prompt provided by the free ChatGPT 3.5 service in December 2023:

[Your Name] [Your Address] [City, State, Zip Code] [Email Address] [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 "Robot-Proof" by Joseph E. Aoun as a highly relevant and engaging resource that aligns with the themes and objectives of my class. As an esteemed publication, I believe this book would greatly contribute to the intellectual exploration of artificial and human intelligences within the context of the novel.

English 347 is designed to provide students with a comprehensive understanding of the intersections between literature and emerging technologies, particularly focusing on the implications of artificial intelligence on human experiences and narratives. "Robot-Proof" appears to be a valuable addition to the curriculum, offering insights into the challenges and opportunities presented by the integration of AI in society.

I kindly request an exam copy of "Robot-Proof" to evaluate its suitability for inclusion in my course. Your assistance in providing this copy would greatly facilitate my course planning process. Additionally, if there are any specific procedures or forms that need to be completed for this request, please do not hesitate to inform me.

Thank you for your attention to this matter. I appreciate your consideration and look forward to the possibility of incorporating "Robot-Proof" into my curriculum. Please feel free to contact me at [Your Email Address] or [Your Phone Number] if you require any further information.

Sincerely,

[Your Name] [Your Position] [Your Institution] 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 applications into teaching can provide educators with powerful tools like personalized learning platforms, intelligent tutoring systems, adaptive learning software, and composition assistants. However, the successful implementation of these technologies requires comprehensive support and training for faculty members. Furthermore, it is crucial to ensure that the human element – the teachers and students – remain central in this digitized education model.

AI applications in education offer an array of benefits. Personalized learning platforms can tailor educational content to meet individual students' needs, thereby improving learning outcomes. Intelligent tutoring systems can provide supplementary instruction to students outside of traditional classroom hours, offering personalized feedback and guidance. Adaptive learning software can modify teaching materials in real time based on a student's performance, paving the way for a more responsive and flexible learning environment. Moreover, composition assistants can help students improve their writing skills, offering real-time corrections, suggestions, and enhancements.

While the potential benefits of AI in education are immense, their successful implementation hinges on providing faculty with adequate support and training. Teachers play a central role in the learning process, and their position cannot be replaced by AI. However, AI can augment their capabilities and make their work more efficient. To leverage these benefits, teachers need to understand how to use these tools effectively and integrate them into their teaching methods.

Training programs tailored to educators should focus on the operation and application of AI tools, as well as strategies for combining traditional teaching methods with these new technologies. Additionally, ongoing support is necessary to help educators troubleshoot issues, stay updated with the latest developments, and learn how to use new features and functionalities as they are rolled out.

The integration of AI into education should be approached with a human-centered perspective. While AI can enhance the learning process, it should not replace the essential human elements of teaching and learning. The relationship between teachers and students, the personal touch in instruction, the understanding and empathy that educators bring to their roles – these elements make the learning process meaningful and effective.

AI tools should be used to empower teachers, not replace them. In a similar vein, they should enhance students' learning experiences, not dictate them. A balance must be maintained where AI tools are used to augment human capabilities and enrich the learning environment, while preserving the humanistic elements of education.

The integration of AI applications into teaching is a promising development in the field of education, offering potential improvements in efficiency, personalization, and responsiveness. However, the successful implementation of these tools requires comprehensive faculty support and training, ensuring educators can effectively leverage these technologies. Most importantly, a human-centered approach must be maintained throughout this process to ensure that the human teachers and students remain at the heart of the learning process. Through this balanced approach,

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 outline that could reveal missing elements.

Our main point in this section is simple: genAI is very far from a panacea that eliminates the need for expert teachers, writers, artists, coders, etc. But leveraged thoughtfully and strategically, with humans still central and responsible, genAI tools can be powerful helpers. Whether in the form of composition assistants like those utilized here or such applications as personalized learning platforms, intelligent tutoring systems, or adaptive learning software, generative AI has capabilities that can meaningfully augment learning processes and workflows for faculty and students alike.

Syllabus policy options and recommendations for helping students use genAI responsibly

In light of the information above, we hope it is obvious that faculty need clear syllabi policies on genAI that reflect their courses' learning goals.

We recommend that our colleagues select among one of three WIU approaches to genAI usage in each of their courses: *permitted*, *partially permitted*, or *not permitted*.

- genAI permitted courses allow this technology's usage in most course assignments and tasks, with exceptions only for rare, explicitly announced contexts like in-class exams. This still does not mean that "anything goes." In such settings, it is crucial that students:
 - commit to ethical genAI use that expands rather than replacing learning;
 - receive respectful attention to their own ethical or privacy concerns, including provision of alternative assignments where possible substitutions could achieve key course aims;
 - fact-check all genAI materials, recognizing the potential for hallucination and inaccuracy;
 - protect the privacy and intellectual rights of all students, faculty, and subjects of study, particularly by avoiding entering sensitive or private information (including FERPA or HIPPA protected information, confidential research information, or copyrighted information) into a chatbot or other genAI interaction context;
 - accept full responsibility for submissions, understanding that utilizing any particular tool does not reduce their accountability for work submitted under their name (including malicious code, harassing statements, or copyright violations);
 - disclose and be prepared to provide complete records of genAI usage (or if faculty prefer, automatically provide such records as appendices, as desired on given assignments).

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

This is a *genAI permitted* course. As such, the use of tools like ChatGPT and DALL-E for course deliverables are welcome, with the exception of in-class quizzes. 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 students, faculty, and subjects; 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.

genAI partially permitted courses allow this technology's usage in some course assignments and tasks, but not others, and carefully defines and announces the parameters for use and non-use. All

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;
- o 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!

- genAI not permitted courses 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 in-class illustrative experiment. In such contexts, we encourage our colleagues to explain their reasons for this choice so as to increase students' active buy-in. Reasonable and appropriate rationale might include:
 - the course's emphasis on student development of entirely original, largely unassisted creative and/or critical thinking;
 - course goals that involve specific skill development, whether for cases in which genAI is unavailable or for more deeply recognizing or internalizing the processes involved; or
 - ethical commitments and priorities that rule out risks and costs involved in contemporary genAI utilization such as those identified in our <u>Opportunities and dangers</u> section.

A genAI not permitted course might include a syllabus statement like the following:

This is a *genAI not permitted* course. This means that genAI tools are completely excluded from our learning environment, except as introduced briefly and minimally by faculty for the sake of demonstration. The reasons for this are not technophobic; there exist worthwhile uses for these tools. However, in this course we will be pursuing aims best achieved without their assistance, including the development of unique writing styles and processes of critical thought. Please understand that in making all course assignment submissions, you are affirming your creation and editing of these documents completely independently of genAI tools (except basic grammar and spellcheck tools built into Google Docs, Microsoft Word, and similar word processors).

In addition to offering these three main categories of genAI usage and non-usage, we wish to make several additional recommendations:

- Because (a) we deeply value the academic freedom of faculty, (b) the use-cases of genAI in our courses is already complicated and diverse, with some faculty very reasonably applying different expectations to different courses; and (c) genAI continues to change rapidly, we do *not* presently recommend a singular AI statement for WIU syllabi. We believe students can benefit more from faculty creating usage and non-usage statements that are more fully tailored to their learning contexts and goals.
- However, we do recommend that all courses should be clearly and officially designated as best fitting one of these three categories (genAI permitted, partially permitted, or not permitted). Furthermore, as indicated above, we urge as much specificity as possible in syllabi and assignment instructions about the uses of particular genAI that are or are not acceptable.
 - The thoughtful "Generative AI" webpages provided by Academic Technology at the University of Michigan suggest the following assignment instruction language, which we

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:

 \rightarrow 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;

 \rightarrow 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;

 \rightarrow you assert that you have respected all specific requirements of your assigned work, in particular requirements for transparency and documentation of process, or have explained yourself where this was not possible.

If any of these assertions are not true, whether by intent or negligence, you have violated your commitment to truth, and possibly other aspects of academic integrity. This constitutes academic misconduct.

- Just as Michigan's language points to the foundational nature of academic integrity, we believe that the new possibilities occasioned by genAI necessitate another revision of WIU's Student Academic Integrity policy. By expanding this umbrella to include usage of genAI according to a given course's category, WIU can help faculty and students to rely on clear shared expectations for learning goals and tools.
- To enable ongoing oversight of AI and digital technology issues, the present ad hoc AI Task Force recommends Faculty Senate strongly consider a new standing committee to focus on these topics. It would be responsible for regularly reviewing genAI and related AI developments, proposing policy updates, and taking other steps to keep faculty and students up-to-date. This committee's purposes could overlap with those of other university technology-related committees, but would prioritize faculty and student learning. The committee could help ensure appropriate ongoing professional development around genAI, whether by updating documents like this one, partnering with other WIU groups to host internal workshops and training sessions, or working to enable additional learning via external resources.
- Finally, we recommend that so long as genAI technology develops in sustainable and affordable directions, WIU should seriously consider university-wide genAI licenses. Currently, faculty and students interested in testing the most advanced varieties of LLM and other genAI must pay monthly subscriptions of roughly twenty dollars per product. Whether one is interested in the unique capabilities of one or multiple tools, this becomes unrealistic for many people quickly, which means that most WIU faculty and students will operate only with free, publicly available versions that are considerably less advanced. It is all too easy to see yet another higher education arms race developing, as many flagship and elite institutions are already forging unique partnerships with AI companies. With any new technology, the danger of rushing forward too eagerly must be weighed against the potential costs of delayed access and learning. GenAI represents a field in which WIU faculty and students could move closer to the cutting edge (both in the technology's use and its critical assessment), should our leaders invest the necessary resources.

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 promises and dangers. Then, we must proceed with measures of caution and boldness that will differ significantly by discipline, learning goals, and individual pedagogies. Whether individually we lean enthusiastic or resistant, we must also stay curious: perhaps the most telling sign that someone should not be trusted about AI is if they claim to have figured it out. Even for people dedicated to this field for their entire career, there are often as many questions as answers.

That said, we can develop better questions, so we conclude this document with what strikes us as some of the most significant questions about AI's trajectory:

- How can we nudge AI's and genAI's impacts on WIU students toward positive changes rather than protection of the status quo? This technology is extremely good at predicting desired (future) information on the basis of the most common (past) inputs, and that is both strength and weakness. One side of the coin is unprecedented ability to find needles in haystacks and discern patterns too subtle or spread across too much data for human eyes to notice. The other side of this coin, though, is genAI's tendency to replicate and buttress already existing attitudes and understanding. GenAI uses the most common patterns in its training data to anticipate the language and shapes most likely to please most users. Upside: it's often right, especially when you're searching for what most people would say. Downside: sometimes most people are wrong or shortsighted, making genAI likely to exacerbate problematic social patterns and assumptions. In order for WIU to be as healthy a learning community as possible, we need to inspire new directions, not wear out old ones.
- If the current builders of genAI are for-profit technology companies with mixed records in prioritizing broad societal good, how can we help keep this technology from exacerbating existing wealth and power gaps or blinding us to its reliance on majority attitudes? Many of our students do not come from privileged backgrounds, so it is all the more important that we work to ensure that students with lower access to and interest in expensive technologies have equal opportunities to learn to marshal AI effectively and ethically. It will be especially important to help them move beyond the equal and opposite mistakes of using genAI simplistically (i.e., uncritically and uncreatively) or of fearing it so much as to avoid learning about it.
- How do we help students want to learn, so that genAI use enhances rather than replaces growth? This is a broader philosophy of education question that may come into increased focus as genAI grows more available and familiar. If our aim is merely to teach regurgitation of information, genAI is already better at that task than most of us will ever be. Even when we memorize effectively, we eventually forget. And assignments that revolve around restating received written or visual information are particularly easy for genAI to complete effectively. This means that we are being challenged to prove our value as educators in new ways: we must be able to motivate students better than algorithms; we must hear them out and share internalized understandings in ways that engage our students' limited attention, specific social contexts, and unique personalities. In designing assignments, we must find ways to deepen and not just score learning. We particularly suggest developing project completion processes that allow for apprenticeship-style monitoring, correction, and encouragement, rather than placing all the weight on the final state of a test or paper.

- What new uses will be found for genAI, and how rapidly will it move us toward Artificial General Intelligence and apparent AI consciousness? In recent months, Michael Cantor observes, genAI tools have begun serving as "jargon demystifiers," "pitiless critics," "robots with feelings," "sous chefs," "whiteboard interpreters," and "speedy summarizers." Some of the claims for their capacities will no doubt continue to be overstated and decontextualized, but the movement is entirely toward more surprises, not fewer. Just as no one fully understands how genAI's generative properties emerged, it seems very possible that additional leaps ahead lie in the near future. While guarding against the assumption that what goes up must always continue to do so (at least with a constant, uninterrupted pace), we must be at least equally cognizant that genAI could very well still be in the early stages of its expansion.
- How do our visions of AI's impacts on the future of work shape the education we seek to give students? As also suggested in the Opportunities and dangers section, the jobs that currently seem most vulnerable to AI replacement are those involving highly repetitive tasks; information recording, organization, retrieval, and basic analysis; machine control and monitoring; and more rudimentary elements of administration. Meanwhile, the varieties of work most likely to remain insulated from AI impacts are those requiring bodily and emotional intelligence or, put another way, dynamic combinations of physical, cognitive, and emotional skills. Keeping such projections in mind, we should be accentuating the parts of our curricula that most develop holistic, self-driven, creative and critical thinkers, while resisting calls to turn students into interchangeable cogs in career wheels that may suddenly be replaced by new AI applications.

For all of its threats, genAI has many potential virtues. Embraced critically and conditionally, it has real potential to drive our society toward much richer work lives for many people. The same is true for academia and WIU specifically: if we continue to uphold our end-goals of academic excellence, educational opportunity, personal growth, and social responsibility and insist on genAI's capacity to serve as means, not end in itself, there remain many welcome paths forward.

Appendix: Extended bibliography

Pedagogy and individual faculty usage decisions

"Adapting Your Course to Artificial Intelligence." University of Kansas's Center for Teaching Excellence. <u>https://cte.ku.edu/adapting-classes-artificial-intelligence-era</u>

Alexander, Jonathan. "Students Have a Right to Write in an AI Era." Inside Higher Ed (November 22, 2023). <u>www.insidehighered.com/opinion/views/2023/11/22/students-have-right-write-ai-era-opinion</u>

Baker, Cal, C.W. Howell, and Fayrah Stylianopoulos. "To Educate Students about AI, Make Them Use It." *Scientific American* (November 14, 2023). https://www.scientificamerican.com/article/to-educate-students-about-ai-make-them-use-it/

Bhuiyan, Johana. "'Is this an appropriate use of AI or not?': teachers say classrooms are now AI testing labs." *The Guardian* (October 31, 2023).

https://www.theguardian.com/technology/2023/oct/31/educators-teachers-ai-learning-classrooms-misuse

Dumin, Laura. "AI in Higher Ed: Using What We Already Know about Good Teaching Practices." *EdSurge* (October 13, 2023). www.edsurge.com/news/2023-10-13-ai-in-higher-ed-using-what-we-already-know-about-good-teaching-practices

Extance, Andy. "ChatGPT has entered the classroom: how LLMs could transform education." *Nature* (November 15, 2023). <u>https://www.nature.com/articles/d41586-023-03507-3</u>

"Generative AI." Academic Technology at University of Michigan. https://academictechnology.umich.edu/instructional-resources/generative-ai

"Introduction to Prompt Engineering." *Microsoft Azure OpenAI Service Documentation* (July 18, 2023). <u>https://learn.microsoft.com/en-us/azure/ai-services/openai/concepts/prompt-engineering</u>

Mollick, Ethan and Lilach Mollick. "Student use cases for AI." *Harvard Business Publishing* (September 25, 2023). <u>https://hbsp.harvard.edu/inspiring-minds/student-use-cases-for-ai</u>

Roberts, Molly. "AI is forcing teachers to confront an existential question." *Washington Post* (December 12, 2023). <u>https://www.washingtonpost.com/opinions/2023/12/12/ai-chatgpt-universities-learning/</u>

"Teaching and Learning in the Age of genAI." Gettysburg College's Johnson Center for Creative Teaching and Learning. <u>https://genai.sites.gettysburg.edu/</u>

Educational strengths, weaknesses, and theoretical considerations

"AI and education: guidance for policy-makers." United Nations Educational, Scientific and Cultural Organization (UNESCO) (2021). ISBN 978-92-3-100447-6. <u>https://doi.org/10.54675/PCSP7350</u>

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 framework for university teaching and learning." *International Journal of Educational Technology in Higher Education* 20.1 (2023). https://educationaltechnologyjournal.springeropen.com/articles/10.1186/s41239-023-00408-3

Chayka, Kyle. "Your A.I. Companion Will Support You No Matter What." *The New Yorker* (November 13, 2023).

https://www.newyorker.com/culture/infinite-scroll/your-ai-companion-will-support-you-no-matter-what

"Ethical AI for Teaching and Learning." *Cornell University Center for Teaching Innovation*. <u>https://teaching.cornell.edu/generative-artificial-intelligence/ethical-ai-teaching-and-learning</u>

Goodlad, Lauren M.E. "Editor's Introduction: Humanities in the Loop." *Critical AI* 1, no. 1-2 (October 1, 2023). <u>https://doi.org/10.1215/2834703X-10734016</u>

Lanier, Jason. "There is No A.I." *The New Yorker* (April 20, 2023). https://www.newyorker.com/science/annals-of-artificial-intelligence/there-is-no-ai

Lawrence, Alistair and Aaron Yaverski. "With AI Writing, Higher Education Should Prioritise Academic Policy and Student Support." *Times Higher Education* (November 30, 2023). www.timeshighereducation.com/campus/ai-writing-higher-education-should-prioritise-academic-policy-a nd-student-support

Ling, Elizabeth. "AI Widening the Gap between Humanities and STEM Worlds." *Harvard Technological Review* (March 26, 2023).

https://harvardtechnologyreview.com/2023/03/26/ai-widening-the-gap-between-humanities-and-stem-wor lds/

Nelson, Alondra. "AI Is Having a Moment—and Policymakers Cannot Squander the Opportunity to Act." *Cap20* (April 11, 2023).

https://www.americanprogress.org/article/ai-is-having-a-moment-and-policymakers-cannot-squander-the-opportunity-to-act/

Obermeyer, Zad, et al. "Dissecting racial bias in an algorithm used to manage the health of populations." *Science* 366 (October 25, 2019). <u>https://www.science.org/doi/10.1126/science.aax2342</u>

Roose, Kevin. "How Schools Can Survive (and Maybe Even Thrive) With A.I. This Fall." *New York Times* (August 24, 2023).

https://www.nytimes.com/2023/08/24/technology/how-schools-can-survive-and-maybe-even-thrive-withai-this-fall.html

Rudolph, Jürgen, Samson Tan, and Shannon Tan. "ChatGPT: Bullshit spewer or the end of traditional assessments in higher education?." *Journal of Applied Learning and Teaching* 6.1 (2023). https://journals.sfu.ca/jalt/index.php/jalt/article/view/689/539

Sarker, Iqbal H. "Machine Learning: Algorithms, Real-World Applications and Research Directions." *SN Computer Science* 2 (2021). <u>https://doi.org/10.1007/s42979-021-00592-x</u>

Somers, James. "A Coder Considers the Waning Days of the Craft." *The New Yorker* (November 13, 2023). <u>https://www.newyorker.com/magazine/2023/11/20/a-coder-considers-the-waning-days-of-the-craft</u>

Tiku, Nitasha, Kevin Schaul, and Szu Yu Chen. "This is how AI image generators see the world." *The Washington Post* (November 1, 2023).

https://www.washingtonpost.com/technology/interactive/2023/ai-generated-images-bias-racism-sexism-st ereotypes/

Weingarten, Randi. "AI is a) powerful, b) perilous or c) both." The American Federation of Teachers (October 15, 2023). <u>https://www.aft.org/column/ai-powerful-b-perilous-or-c-both</u>

Broader societal and ethical implications and new/future contexts

Chiang, Ted. "Will A.I. Become the New McKinsey?" *The New Yorker* (May 4, 2023). <u>https://www.newyorker.com/science/annals-of-artificial-intelligence/will-ai-become-the-new-mckinsey</u>

Dimock, Wai Chee. "AI and the Humanities." *Proceedings of the Modern Language Association* 135 (2020). <u>https://doi.org/10.1632/pmla.2020.135.3.449</u>

Dzieza, Josh. "AI Is a Lot of Work." *Intelligencer* (June 20, 2023). https://nymag.com/intelligencer/article/ai-artificial-intelligence-humans-technology-business-factory.html

Edwards, Benj. "Mother plucker: Steel fingers guided by AI pluck weeds rapidly and autonomously." *Ars Technica* (November 28, 2023).

https://arstechnica.com/information-technology/2023/11/mother-plucker-steel-fingers-guided-by-ai-pluck-weeds-rapidly-and-autonomously/

Etienne, Mill. "Despite OpenAI Shakeup, Augmented Reality Is Game-Changer, Neurosurgeons Say." *Forbes* (November 25, 2023).

https://www.forbes.com/sites/milletienne/2023/11/25/despite-openai-shakeup-augmented-reality-is-gamechanger-says-neurosurgeons/?sh=67be23d93cd6

Evaristo, Bernardine. "It is a beast that needs to be tamed': leading novelists on how AI could rewrite the future." *The Guardian* (November 11, 2023).

https://www.theguardian.com/books/2023/nov/11/chatgpt-needs-to-be-tamed-bernardine-evaristo-jeanette -winterson-put-ai-to-the-test

Friis, Simon and James Riley. "Eliminating Algorithmic Bias Is Just the Beginning of Equitable AI." *Harvard Business Review* (September 29, 2003). https://hbr.org/2023/09/eliminating-algorithmic-bias-is-just-the-beginning-of-equitable-ai

Foy, Kylie. "New tools are available to help reduce the energy that AI models devour." *MIT News* (October 5, 2023).

https://news.mit.edu/2023/new-tools-available-reduce-energy-that-ai-models-devour-1005

Harper, Kelso and Sophie Bushwick. "Talking to Animals with Artificial Intelligence." *Scientific American* (December 27, 2023).

https://www.scientificamerican.com/podcast/episode/talking-to-animals-with-artificial-intelligence/

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* (December 22, 2023). <u>https://www.nature.com/articles/d41586-023-04094-z</u>

Kochhar, Rakesh. "Methodology for O*net Analysis." Pew Research Center's Social & Demographic Trends Project (July 26, 2023).

https://www.pewresearch.org/social-trends/2023/07/26/2023-ai-and-jobs-methodology-for-onet-analysis/

Kozlov, Max. "AI that reads brain scans shows promise for finding Alzheimer's genes." *Nature* (November 10, 2023). <u>https://www.nature.com/articles/d41586-023-03482-9</u>

Krügel, Sebastian, Andreas Ostermaier, and Matthias Uhl. "ChatGPT's inconsistent moral advice influences users' judgment." *Nature* (April 6, 2023). <u>https://www.nature.com/articles/s41598-023-31341-0</u>

Lozlov, Max. "AI that reads brain scans shows promise for finding Alzheimer's genes." *Nature* (November 10, 2023). <u>https://www.nature.com/articles/d41586-023-03482-9</u>

Masterson, Victoria. "AI can now outperform conventional weather forecasting – in under a minute, too." *World Economic Forum* (December 14, 2023). https://www.weforum.org/agenda/2023/12/ai-weather-forecasting-climate-crisis/

Mims, Christopher. "AI Is Ravenous for Energy. Can It Be Satisfied?" *Wall Street Journal* (December 15, 2023). <u>https://www.wsj.com/tech/ai/ai-energy-consumption-fc79d94f</u>

Perrigo, Billy. "OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic." *Time* (January 18, 2023). <u>https://time.com/6247678/openai-chatgpt-kenya-workers/</u>

Savage, Neil. "How robots can learn to follow a moral code." *Nature* (October 26, 2023). https://www.nature.com/articles/d41586-023-03258-1

Schmidt, Eric. "Eric Schmidt has a 6-point plan for fighting election misinformation." *MIT Technology Review* (December 15, 2023).

https://www.technologyreview.com/2023/12/15/1085441/eric-schmidt-plan-for-fighting-election-misinfor mation/

Singh, Maanvi. "As the AI industry booms, what toll will it take on the environment?" *The Guardian* (June 8, 2023).

https://www.theguardian.com/technology/2023/jun/08/artificial-intelligence-industry-boom-environment-t oll

Wells, Sarah. "Generative AI's Energy Problem Today Is Foundational." *IIEE Spectrum* (October 29, 2023). <u>https://spectrum.ieee.org/ai-energy-consumption</u>

Verma, Pranshu. "AI fake nudes are booming. It's ruining real teens' lives." *Washington Post* (November 5, 2023).

https://www.washingtonpost.com/technology/2023/11/05/ai-deepfake-porn-teens-women-impact/