Talking about Generative AI
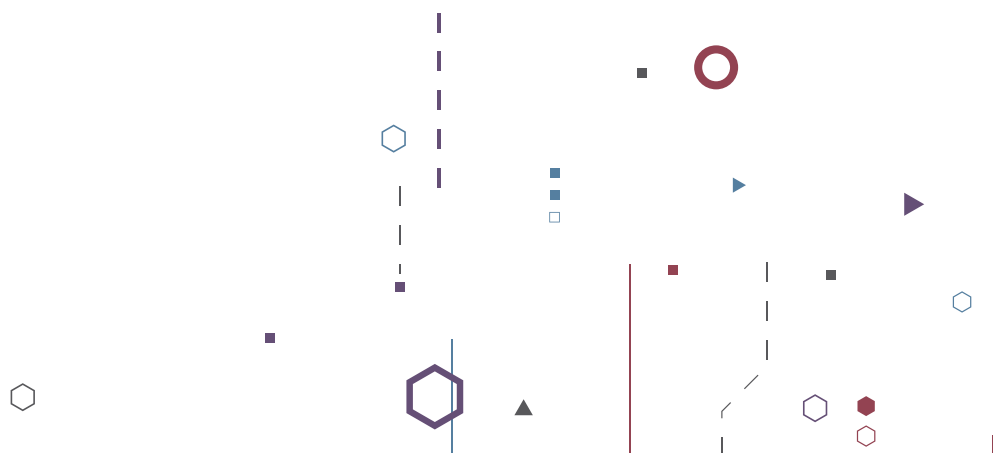
A GUIDE FOR EDUCATORS

Sidney I. Dobrin

Version 1.0

Contents

Goals of This Resource

The introduction of ChatGPT has been a wakeup call, and as often happens when an alarm sounds, some people are startled and frightened, some are angered, some are slow to react, some are ignoring it altogether, and some are embracing it. ChatGPT, and Generative AI (GenAI) more broadly, signals a significant change in how we think about our roles as educators and in how we assess the function of higher education.

You've probably heard the now-clichéd reaction that AI is going to take away people's jobs—to render them obsolete in the work world. The truth is that people will not lose their jobs to AI; they will lose their jobs to people who know how to work with AI. Those who understand how to interact with AI will thrive beyond those who don't. This is, in part, why educational institutions need to address the role higher education will play in preparing students for career success in an AI-driven workforce.

Talking about Generative Artificial Intelligence is designed not to provide answers to the questions that have emerged alongside and in response to the GenAI revolution, but rather to provide points of departure for discussion, thought, and research, both locally and globally. This resource is provided to assist administrators, instructors, and researchers in engaging with GenAI in the ways that best serve their campuses, their classrooms, their research, and their students. It is a guide to joining the conversation.
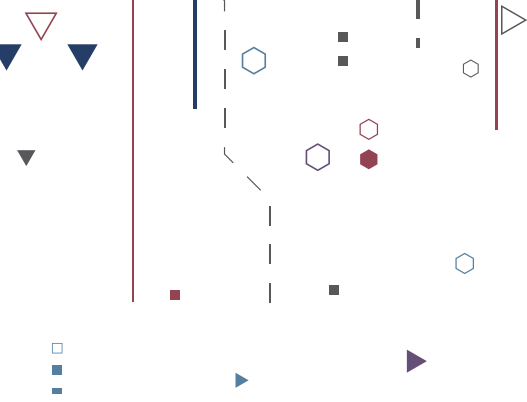
01

What Is Generative Artificial Intelligence?

What Is Artificial Intelligence and How Does It Work?

In this resource, I'm not going to provide a comprehensive history of artificial intelligence (there are several great books out there that do this; see, for example, Michael Wooldridge's *A Brief History of Artificial Intelligence*), nor am I going to unpack all of the nuances of the various functions of AI. However, it's important to acknowledge that Artificial Intelligence emerged as an umbrella term in the 1950s to describe hypotheses about, and attempts to build, computers that could perform tasks in the same ways the human brain does. The idea was to make a computer that could complete functions that previously only human brains could accomplish, including things like complex problem solving, communication through natural language, reasoning, and visual recognition. The AI machine, like a human, would "learn" from its experiences with these tasks, making it more adept and efficient in future performance.

Thus, we now familiarly—and reductively—define **Artificial Intelligence** as both the theory and development of computer systems that can perform tasks that previously required human intelligence.

This definition identifies two key approaches to how we think about AI: the theoretical, or what we might call the conceptual; and the applicable. Thus, we can think about AI from both perspectives: Conceptual AI and Applied AI.

Conceptual AI refers to how we think about the ramifications of AI. It asks questions about how AI will impact societies, economies, and cultures. Conceptual AI stands at the center of the questions we face as educators, especially regarding how we teach and

how students learn. Conceptual AI takes up the main ethical issues that surround AI's evolution and use. Conceptual AI also theorizes about how we might further develop specific AI technologies. It focuses on questions of *Why?* or *What if?*

Applied AI asks questions about how to use AI, both generically and specifically. It focuses on practicality and use. Applied AI is the approach we teach students for developing a facial recognition system, or for using GenAI to write an essay on *Moby Dick*. Applied AI focuses on questions of *How?*

Conceptual AI and Applied AI are two sides of a single coin, as both are necessary for understanding, making, and using AI technologies. To engage AI responsibly, administrators, instructors, and students will need to understand the hows, the what ifs, and the whys of AI.

Part of the reason we should consider both the conceptual and applicable aspects of AI is that AI operates differently from traditional computer programs, both structurally and functionally. Computer programs are, in essence, sets of linear instructions delivered to the computer. For example, one of the first programming strings coders using Basic are taught is the If/Then structure. It's a basic formula used to convey "if this, then that." In code, it consists of a straightforward, linear set of instructions looking something like this:

```
001 If A = 10 then goto 025
002 If A < 10 then goto 026
003 If A > 10 then goto 027
```

However, AI doesn't operate in this linear fashion. Instead, AI uses **algorithms** that "learn" from each engagement and then make predictions based on previous experiences in performing the same or a similar task. This is known as **machine learning**. Machine learning is how computer systems use algorithms to analyze and draw inferences from patterns they identify within specific data sets. That is, when an AI

recognizes patterns within a data set, it "learns" to make inferences about those patterns. Machine learning requires AI systems to "train" using predetermined data sets in order to "learn" to identify patterns in those sets. Machine learning is how a computer learns from the data it encounters in order to perform tasks that it hasn't been specifically programmed to perform. This is a central component of GenAI.

In addition to machine learning, AI algorithms can also use **deep learning** techniques for processing information. Deep learning can be thought of as a series of complex algorithms modeled on the human brain and its structures. This learning technique allows the AI to work with varied data, such as images, text, and sound.

## What Is Generative Artificial Intelligence and How Does It Work?

Researchers made only minimal advances in developing AI between the early explorations of the 1950s and the early 2000s. Then, in 2012, the idea of the **neural network** was developed. A neural network is a mathematical system that scans large amounts of data to identify patterns. The AI processes the data in those patterns to "learn" the characteristics of the data in the pattern. For example, an AI might identify a recurring pattern showing that *Moby-Dick* was a novel written by Herman Melville in 1851. It might also discern from those patterns that *Moby-Dick* is the title of the book and 'Moby Dick' the name of the whale. It can only do so, however, if there are multiple statements identifying this information, sufficient to create a pattern. Similarly, a neural network can look for patterns in images of a whale, in order to identify what a whale is. This is how voice recognition AI is able to understand what you mean when you say "pay bill" in response to an automated customer service prompt on the telephone.

In 2014, computer scientists developed the idea of **generative adversarial networks** (GANs), a variation on a machine learning algorithm that allows AI neural networks to take the data they gather through machine learning and use it to generate what appear to be new, or even original, iterations of that information through written text, images, and sounds. GANs provided a way for AI platforms to take what they have "learned" and produce new texts based on that information. Early GAN models were restricted in their data sets, limiting the extent of what they could generate.

Around 2018, several AI-focused companies, including Google and Microsoft, began building and using neural networks that could scour and scrape data from the massive and complex information available on the Internet. These new neural networks could access data from Wikipedia, from databases of digitized books, from academic publications, from social media, and from nearly anything else available on the Internet. These extensive data sets are known as **Large Language Models** (LLMs). By identifying patterns within LLMs and then reorganizing the information in those patterns, AI machines become able to reproduce those patterns in their output, which would appear to be original and in the style of a human writer or image maker. This is **Generative Artificial**

Intelligence—or "GenAI," the primary subject of this resource, and the category of AI that ChatGPT falls under.

GenAI, then, is a class of AI algorithms that can create a variety of content types—including alphabetic text, images, and sound—based on patterns it has identified within a data set. The algorithm has "learned" to identify specific characteristics within that pattern.

Prompt

↓

Scrub

↓

Output
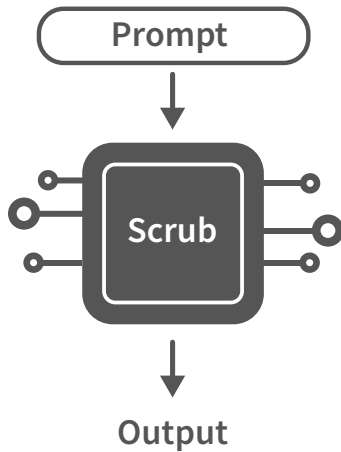
GenAI works like this: a user provides the AI with a prompt asking the AI to create a specific deliverable—an essay, a song, an image, the solution to a math problem, or so on. The AI then scrubs through all of the data available, looking for patterns and recurring information about the requested task. It then reorganizes that data into a pattern that it deems to answer the prompt. (This is a very simplified summary of the process.)

We can think of a GenAI as participating in a rudimentary conversation with a user. The user asks a question, then the AI locates information and attempts an answer. This is how most ChatBots function—hence the name "ChatGPT" (Chat Generative Pre-trained Transformer). ChatGPT is just a more complex chatbot, pre-trained to locate data, transform that data, and generate new ways of conveying the data.

▶ **THE CONCERN ABOUT** GenAI's restriction to digital data ties directly to the ways in which we teach students about research methodologies. However, this issue is not unique to GenAI. Over the past decade, several graduate and undergraduate students have told me that "if it's not online, it's not usable research." This emerging research zeitgeist speaks to several issues higher education will need to continue to take up, including the ways in which we account for non-digitized research availability, the function of non-digitized texts, and even the evolving understanding of what a campus library provides. Limiting data to what is accessible online leaves a tremendous amount of research outside of the "useful" pool. Such limits also bring forth important questions of inclusion and representation (as discussed in Part III of this resource). And they bring to bear the distinctions we make about what counts as digitization, as simply scanning documents doesn't necessarily make them accessible if they cannot be scrubbed for individual terms, images, and such.

The introduction of GenAI technologies requires that we consider the availability of digital research in two key ways: a) in the accessibility of research to the LLMs, and b) in how we negotiate student understanding of what constitutes viable research. (Yes, it also raises the question as to the role of library facilities, but that's a different type of institutional question.)

There are several weaknesses, however, in current iterations of GenAI. For one, current AIs are not able to discern value or accuracy or bias in the data they scrub. They are only able to identify patterns within the data, and thus they often incorporate inaccurate, false, or biased information. Current GenAI programs are also unable to identify which data is more or less relevant in light of the provided prompt. **For these reasons, it is essential to teach students the critical skills of assessing accuracy and relevance if they are to use these GenAIs effectively.**

GenAI systems can also only gather information from the LLMs to which they have access. So, for example, research or data that hasn't been digitized will never be part of what the AI can draw on. This limit restricts what an AI can generate, and may bias results in favor of data this is widely and digitally available.

Generative Artificial Intelligence and Post-Secondary Education

Perhaps you've heard that AI in general, and GenAI in particular, is destroying education—that it allows high school and college students to easily cheat on their essay assignments or their computer programs or their mathematical proofs. Many are worried that it allows students to pass MBA or MCAT exams, or to complete dozens of other tasks that instructors have traditionally assigned in order to assess student mastery and knowledge. Perhaps you've seen the calls to ban students from using such GenAI entirely. But perhaps, alongside these arguments, you've also seen claims that GenAI is revolutionizing pedagogy and pointing us to a new paradise in higher education.

No matter what you've heard, the fact of the matter is that GenAI is one of the most ground-shaking technological advances that higher education has encountered. GenAI's emergence and evolution have unfolded with such speed that higher education is just beginning to explore the relationships between GenAI and teaching, learning, and research.

To understand just how rapidly GenAI has advanced, consider that ChatGPT was only publicly launched in November 2022. Within five days, it logged in over a million users. No other application has achieved that size of user base in that short of a time. And while it took TikTok, the popular social media platform, nine months to reach 100 million monthly users, it took ChatGPT only two.

Part of ChatGPT's rise to fame stems from its versatile application to a broad range of use cases. It can create long-form written content such as essays; it can write computer programs; it can score highly on the MCAT; and it can pass an MBA licensing exam. ChatGPT's release also signaled a shift in the availability of GenAI, moving it into the hands of everyday people rather than just the AI experts.

As the barrage of media coverage about GenAI and higher education shows, colleges and universities are having difficulty finding ways to address its use among students. While a handful of institutions have already created formal policies, most have

TIME TO 1 MILLION USERS

ChatGPT — **5 days**
Spotify — 5 months
Facebook — 10 months
Twitter — 2 years
Netflix — 3.5 years
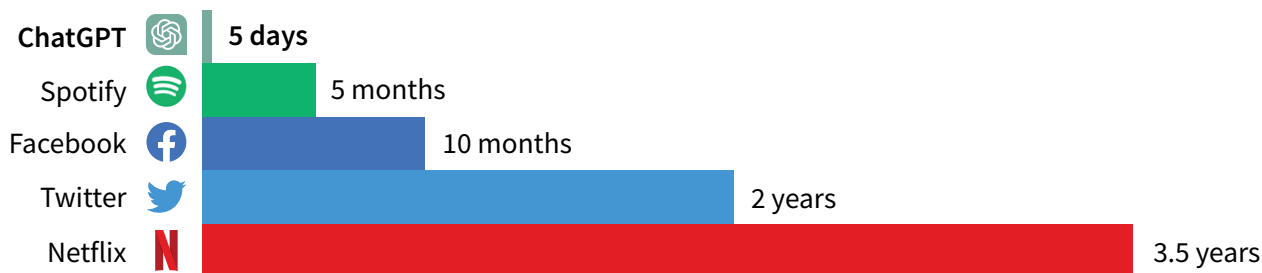
Source: Buchholz, Katharina, and Felix Richter. "Infographic: ChatGPT Sprints to One Million Users." *Statista Infographics*, 24 Jan. 2023.

TIME TO 100 MILLION USERS

ChatGPT — **2 months**
TikTok — 9 months
Instagram — 30 months
Pinterest — 41 months
Spotify — 55 months
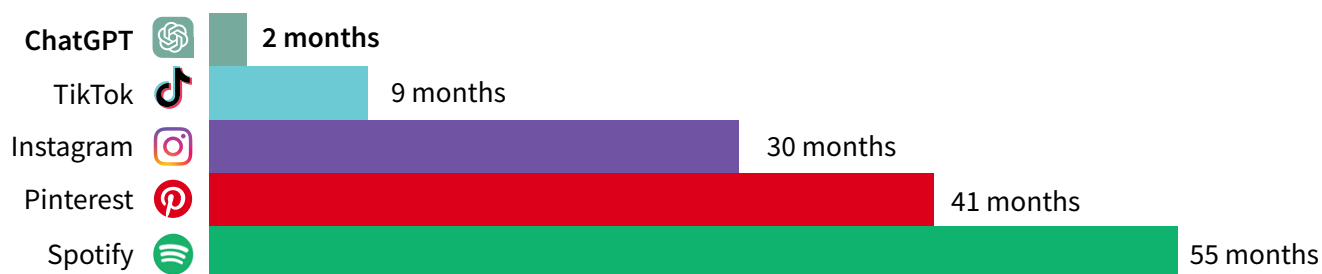
Source: "CHATGPT's on Track to Surpass 100 Million Monthly Users Faster than TikTok or Instagram: UBS." *Yahoo! News,* Yahoo!.

not, and many acknowledge that they are not prepared to offer directions to either instructors or students regarding the responsible use of GenAI. In a revealing survey **published by** *Best Colleges* in March 2023, 54% of college students surveyed claim their instructors have not talked with them about GenAI tools such as ChatGPT. Tellingly, the same survey reveals that 61% of students surveyed believe GenAI will become part of the "new normal" in higher education.

History of Writing Technologies and Resistance to Them

Imagine, if you would, a writing class that bans word processors, spell checkers, grammar checkers, or citation generators—technologies that, when they first appeared, some teachers vilified as the death of writing education. Or consider the moral outrage levied against the electric calculator in the 1970s and the debates about its effect on student learning. Outrage against emerging and evolving technologies that affect student learning—or, perhaps more accurately, that disrupt the comfort of our familiar pedagogies—is certainly not new.

All technologies—not just digital technologies or writing technologies, but *all* technologies—have two paths: either they become ubiquitous and naturalized into how we do things, or they become obsolete. Sometimes they take both paths. In most cases, they

become obsolete because another technology has surpassed the old technology's usefulness or efficiency. In others, they become ubiquitous because they serve the needed purpose well enough—at least for a time. No one picks up a spoon and marvels at it as a magnificent piece of technology. It's a ubiquitous technology, but it's a technology that significantly changed human culture. It altered social dynamics; it altered public health; it affected art and aesthetics; it impacted labor practices; it encouraged other technological developments. Technologically speaking, spoons and writing are quite similar. They are ancient technologies that have become ubiquitous, have had significant impacts on the world, and are rarely thought of as technologies. (Granted, writing is undoubtedly responsible for causing more deaths throughout human history than spoons have, but you get the point of the comparison.)

Media theorist Gregory Ulmer explains that when human cultures evolved from oral to literate, fundamental aspects of our lives changed (some for better, some for worse). The same change is unfolding now as we move from a literate culture—a culture based in print reading, writing, and communication—to a post-literate culture, or a digitally literate culture, which Ulmer calls "electracy." Like the shifts from oral culture to literate culture, the shift to digital culture will inevitably change much about us as human beings. These are big techno-philosophic questions not taken up in this resource. However, the uncertainties raised by these big questions may explain some of our common worries regarding GenAI: that education won't retain the processes and systems we've come to accept as normal; that it will disrupt practices we've relied on for thousands of years; and that it will force changes on the many of us who've invested our careers in educational practice. GenAI may seem risky to us because it makes us rethink the things in which we are comfortably entrenched; it asks us—rather abruptly, it might seem—to change.

Generative AI is not destroying education; it is reinvigorating education in a new context. It is rewriting it. And that requires us as researchers, administrators, and instructors to think about some of the fundamental values and practices on which we've come to rely.

That doesn't mean we have to abandon everything we know about education, or about good teaching, or about what it means to learn. It means we have to integrate the old and the new. It means we have to apply some of our tried-and-true approaches in these new contexts, and we have to change or even abandon some other ones.

Generative Artificial Intelligence Is Not as New as You Think

In December 2013, the *Los Angeles Times* published a short article, excerpted here:

> A shallow magnitude 4.7 earthquake was reported Monday morning five miles from Westwood, California, according to the U.S. Geological Survey. The temblor occurred at 6:25 a.m. Pacific time at a depth of 5.0 miles....
>
> *Los Angeles Times*, December 23, 2013

Though many of us had been unaware of the use of GenAI bot writers until the recent media attention, AI writers have been churning out content for at least a decade in places we might not expect. The *LA Times* article quoted above was written by a bot known as "QuakeBot." The newspaper employs QuakeBot to write news articles about earthquakes. Connected to U.S. Geological Survey monitoring and reporting equipment, QuakeBot can produce an article nearly instantly, containing all of the relevant—and accurate—information readers need to have: where the earthquake was centered, the magnitude of the quake, aftershock information, and so on. The article is then transmitted to readers, who are unlikely to question whether the article was bot- or human-written, because the author is likely irrelevant in the conveyance of this information. QuakeBot now tweets information to readers under the Twitter handle "LA QuakeBot".

AI writers are much more ubiquitous than most of us recognize. The international news agency *Bloomberg* has for years relied on automated writing technologies to produce approximately one third of its published content. The *Associated Press* uses GenAI to write stories, as does *The Washington Post*. *Forbes* has tried using GenAI to provide reporters with templates for their stories for years now.

Journalism is an interesting example, as it's a domain in which many assume that human writers/reporters do all of the work of research and writing. It's also a communicative arena in which the idea of *integrity* is central (more on this in Part II). Journalism is an area in which GenAI technologies have been assimilated, and have contributed to the understanding of what it means to conduct research and to produce writing—the very questions higher education must now engage across all disciplines.

Beyond journalism, we've been interacting with AI technologies for some time now, through technologies such as chatbots and phonebots. AI-generated content has become omnipresent in our lives, even if we are often ignorant of this fact.

WHEN WRITING the word "phonebots" at the end of this section, the spell checker (a type of GenAI technology), wanted to change "phonebots" to "phonebooths," a cultural reference to a nearly obsolete technology, one with which it's unlikely our students are familiar. Even the GenAI technology can't keep up with the ways in which our language and cultural references embody the use of our technologies. That moment brought to mind, too, the ways in which our slow-to-change cultural attachments to technology are made evident in our discourse. Once, I told my son to "hang up the phone" and he had no idea what that meant, since the expression has become obsolete.

02

Artificial Intelligence for Course Instructors

Generative AI and Integrity

Upon ChatGPT's release, as popular media became flooded with stories about the future of education in the era of GenAI, a significant conversation began around matters of academic integrity. Specifically, many claimed that ChatGPT and other GenAI platforms would facilitate widespread cheating, plagiarism, and other forms of academic dishonesty.

The prevalence of these discussions foregrounds two worrisome assumptions about secondary and higher education: a) that students will always cheat if given the means, and b) that a primary function of instructors and institutions of higher education is to police students and their use of such tools. These attitudes about students and college education are mostly inaccurate.

The question of how we can prevent students from using GenAI to cheat is a question of surveillance and policing, not of education. It is not the question educators should be asking.

Instead, educators should be asking how we can teach students about the relationships between GenAI and integrity—academic, professional, civic, and personal. These are the kinds of questions that facilitate teaching and learning; these are the kinds of questions that drive dialogue, collaboration, inquiry, and curiosity. If the conversation about GenAI starts with questions about students' own integrity, then the conversation evolves within an atmosphere of negativity and distrust rather than in a climate of encouragement and possibility.

We can think of academic integrity as the ethical code of conduct for higher education. Academic integrity takes up matters of academic honesty, cheating, and plagiarism for students. It also addresses honesty and rigor in academic research and publication. Fairness, responsibility, and trust are its core values. There are also powerful and important connections between academic integrity and matters of diversity, equality, and inclusion.

Academic integrity is certainly important, and there's no aspect of it that doesn't have direct and immediate relevance to discussions about GenAI in education. However, higher education should take up GenAI not just in relation to *academic* integrity but also to professional, civic, and personal integrity, particularly as we ask students to think about how what they learn in our institutions will transfer to their lives. Whenever we teach students, part of our assumed mission is to prepare them for the things they will do when they leave our guidance. As such, conversations about academic integrity and GenAI need to cover more than just plagiarism, original work, and authorship. **We need to understand that students will inevitably interact with GenAI technologies in just about every aspect of their lives.**

For example, it would be of benefit to students—and, frankly, to instructors as we learn about GenAI—to ask them to investigate the ways in which GenAI technologies are deployed in discipline-specific contexts. How are they used, for example, in medical research, or in the social sciences, or in the humanities? There are also benefits to asking students to explore the ways in which GenAI systems are used in workplace settings. These kinds of inquiries can be situated in terms of integrity while simultaneously serving experiential learning objectives.

The Question of Authenticity

We assess students in classrooms to verify mastery and expression, and to ensure—we tell ourselves—that students have learned content and methods, and that they are able to communicate or express their understandings of those content and methods. Our traditional methods for such assessments revolve around the assumption that we *can* measure and evaluate a student's mastery and expression. However, such measurements are often problematic from the outset.

One of the core questions about students' use of GenAI revolves around the *authenticity* of the work they submit for evaluation. That is, GenAI's emergence makes evident the challenge of instructors' abilities to identify a student's work as actually their own. Much of this quandary is bound up in academic traditions of assessment and the belief that the work a student submits reflects accurately what the autonomous student agent does or does not know. Now, we must ask whether instructors are able to discern between GenAI-produced responses and those written by a student. This expands and complicates the familiar question, "Is the work your own?"

Perhaps more relevant than simply identifying whether a student deploys GenAI in responding to an assignment, though, is the need to identify what part of a student's work is "authentically" the student's and which is GenAI-generated. This makes evident, too, the emerging question of how to assess human-machine collaboration. Collaborative assignments have long caused problems for instructors trying to assess group work (*which students did which parts of this group assignment?*); human-machine collaboration poses a set of questions all its own.

The Plagiarism Problem

As noted, one of the primary conversations in the GenAI orbit has been fear about students using GenAI platforms to cheat on assignments. The *Best Colleges* survey noted earlier found that 51% of students see using GenAI technologies as a form of plagiarism, though about 20% of the surveyed students use those tools anyway, as of March 2023.

Educators around the world have been stymied and left scrambling to contend with the possibilities of students using GenAI tech to produce work that they then represent as their own. It's currently difficult to identify when content has been produced by a GenAI system such as ChatGPT, and there are now resources available to help students mask their use of GenAI. However, plagiarism checkers such as TurnItIn, Grammarly, and Chegg are rapidly developing and integrating GenAI checkers into their systems. At some point, checking for GenAI-generated content will be standard for most plagiarism checking systems, though they are not fully reliable in this capacity yet.

Until then, we need to recognize that identifying GenAI content can be difficult and time consuming. There is no simple answer here. Whether or not a technological solution arises at some point in future, the rapid advancement of GenAI technologies means that skill and adaptation will be required of instructors who wish to reliably identify GenAI content. Adding this degree of scrutiny toward student work can add significant labor for instructors, and extend grading turnaround times. This increase in teacherly responsibilities must be addressed by individual institutions. Campuses need to account for this added labor. Given that first-year writing classes in particular are most often taught by non-tenured faculty or graduate teaching assistants, there is a compelling need to establish best practices for assessing student writing in light of GenAI.

GenAI has clearly outpaced higher education's longstanding (perhaps *stagnant* or *archaic* are better words) understandings of plagiarism and academic integrity. At stake is our very understanding of what counts as plagiarism. That is, the seemingly sudden appearance of accessible GenAI technologies compels higher education to rethink many of its fundamental assumptions and beliefs, including what we mean by *plagiarism*. Traditionally, we understood plagiarism to mean using the words or ideas of another person without accrediting that work to the original writer. This understanding has been so central to the core values of higher education that it spawned detailed

methods and expectations for research and source citation. There would be no need for APA, MLA, Chicago, Turabian, IEEE, and other citation styles without the elemental doctrines of plagiarism and academic integrity. However, the emergence of GenAI calls into question those familiar principles in three critical ways:

▷ First, the customary understanding holds that plagiarism occurs when a person steals ideas or words from another person. However, because GenAI generates data that is not attributable to a person—that is, its output is written by a machine, not a person—we may wonder whether the use of GenAI content falls under the current definitions of plagiarism to begin with. Note, too, that this also raises questions about intellectual property and ownership more broadly, which are bound up in ideologies of capital and property.

▷ Second, the quandaries surrounding plagiarism and GenAI technologies call into question our longstanding acceptance that humans operate as individual writing agents. Now that GenAI technologies can "write" independently of the human, will we be willing to grant them authorial agency over the texts they produce?

▷ Third, because GenAI provides content based upon what it can gather from its LLM, which comprises data that humans and machines have written and which now circulates digitally, we must ask whether some GenAI platforms are themselves guilty of plagiarism, regardless of whether a student later submits the GenAI's output under their own name.

Questions about GenAI and plagiarism grow from issues in both conceptual and applied AI. How do we understand and teach students to use GenAI within the frame of our current definitions of plagiarism, and, conceptually, how will we adjust our understanding of plagiarism in light of GenAI?

(It's important to note as well that concerns about GenAI and plagiarism are very much Western worries, as the very idea of plagiarism is not inherently universal. Therefore, the relationships between institutional definitions of plagiarism and GenAI should also be considered from multicultural/transcultural positions.)

For now, instructors appear comfortable identifying text copied from a GenAI as being equivalent to copying from another student or other source. This approach allows instructors to retain their current understandings of plagiarism, and to keep in place their familiar methods for teaching and policing acts of plagiarism. Longer term, however, **instructors, researchers, and administrators will need to reassess and redefine plagiarism and academic honesty both for student work and for their own research** as GenAI technologies become more specialized and more integrated into our processes for conducting research and producing text.

▶ **JUST AS EDUCATIONAL** institutions will need to rethink plagiarism and academic integrity in the shadow of GenAI, industry will need to address similar concerns. Consider, for example, a magazine or blog in the recreational industry that uses a GenAI platform to write an article about advances in the thermoplastic elastomers used to make hiking shoes. Assume that the topic of the article is timely and important to the industry, and that competing publications have generated similar articles using the same GenAI application. Since the GenAI has scrubbed its data from the data in its accessible LLM, the articles in each publication will inevitably be remarkably similar in the content they convey and, perhaps, in their writing style. How will industry develop its sense of professional integrity surrounding such scenarios? Moreover, how will the "writers" who use GenAI to produce such content be trained in doing so? Or, held accountable for doing so?

Given the relatively recent emergence of concerns about GenAI and academic integrity, we should be cautious, too, of students' understanding of when and how to use GenAI-generated content. That is, we should be alert to students' intent in how they use GenAI. For example, a student who cuts and pastes information from a web source without using citation may have a different intention than a student who uses a GenAI to produce a longer piece of text as a response to an assignment. It's important to recognize that students may not understand the ethics of or policies about using GenAI content, simply because the opportunity to use these platforms is so new and they have not yet learned how to use them ethically. **Talk with your students about how, when, and why they can use GenAI-produced content.**

Designing Assignments that Outpace GenAI?

Concurrent with the ongoing considerations about GenAI and plagiarism, there is much discussion lately about ways to rethink classroom assignments—particularly writing assignments—to make them immune, or at least resistant, to GenAI responses. For many years now, instructors have upheld the value of designing writing assignments that make plagiarism difficult. I have regularly told writing instructors that if students plagiarize responses to an assignment, it's because the assignment design asks for responses that *can be* plagiarized. That is, often the opportunity for plagiarism is as much the fault of the instructor for designing assignments that can be plagiarized as it is of the student who consciously decides to plagiarize.

Some of the strategies that instructors have traditionally espoused as ways to deter students from plagiarizing include:

▷ assignments that specify the use of particular research resources

▷ assignments that require the use of direct quotations

▷ assignments that are so uniquely local to the class that reusable responses wouldn't be available outside of the classroom

- ▷ assignments that require submissions of early and subsequent drafts throughout the writing process

- ▷ assignments that address recent events or new texts about which there isn't yet much accessible information on the Internet

- ▷ assignments that don't rely on traditional essayistic forms as the student deliverable, such that there is no place for students to include plagiarized text

- ▷ assignments that rely on analysis of non-written texts such as images or films

- ▷ assignments grounded in personal experiences, such as literacy narratives or reflection responses about a student's own writing process

When ChatGPT jumped to the fore, this advice re-emerged as a familiar and convenient method retrofitted to address students using GenAI to complete assignments. Many assumed that we could defend against GenAI-generated content by simply doing what we have always done to ward off plagiarism. Early responses to ChatGPT's emergence included recommendations to create distinct assignments on narrow topics, to which the broad, generic responses of GenAI couldn't possibly provide correct or accurate responses. Outsmart the bot. Identify and exploit GenAI's gaps and flaws.

These approaches, however, are often ineffective against GenAI, as its range of skills grows rapidly and often in ways unaffected by traditional methods of plagiarism prevention. Some proposed upgrades to these approaches include requiring citations or requiring academic research sources. Yet, even these additions will not lead to assignments that are immune to GenAI, as more sophisticated applications continue to be released, such as ChatGPT4 and other GenAI programs tailored to specific tasks. Furthermore, each GenAI is perpetually learning from each task it performs, and improving through this process. As Anna R. Mills puts it, "We can't out-prompt it."

Since the rapid evolution of GenAI outpaces instructors' abilities to identify gaps or blind spots to use as student assignments, **it may not be in the best interest of instructors to devote energy toward outpacing GenAI.** That is, trying to develop assignments for which GenAI platforms cannot provide viable responses may be impractical—if not impossible—given the velocity of AI evolution.

Attempts to develop assignments to outpace or outmaneuver GenAI technologies may also reinforce learning objectives that may now be outmoded in the AI era. Given that our familiar approaches to staving off plagiarism will not translate wholesale into measures for preventing students from using GenAI, we need to rethink how we understand the relationship between assignments and GenAI-produced responses. More importantly, we should seriously reconsider our motivations in assignment design if AI-proofing or plagiarism-proofing our assignments is one of our primary objectives. We should be more concerned about what students will *learn* from an assignment and

what *educational purpose* the assignment serves, than whether the assignment's imperfections provide students with opportunities to circumvent the learning objectives. **It may be of greater value for students to learn how GenAI can and cannot inform the ways in which they respond to assignments.** This will require many of us to revise—if not altogether surrender—a lot of the assumptions and philosophies upon which we have developed our teaching careers.

Assignment Design in the Here and Now

Nonetheless, instructors are eager for guidance about accounting for GenAI in the classroom *right now*—not after the prolonged period of academic soul-searching suggested above. Because their students have access to and are using GenAI right now, they need advice right now. Below are some initial ways to think about designing assignments with GenAI in mind:

1. Incorporate more on-demand, in-class writing assignments that require students to engage material with a greater degree of immediacy than a take-home or long-term assignment might require. For example, if a class lecture or discussion has conveyed information about subject X, use part of the class time to have students generate content about subject X, whether that content be a math formula, a short piece of writing, or a list of key terms.

2. Require students to include materials that are only available in your classroom lessons, lectures, or lab work. Such materials are not likely to be available in the GenAI's LLM, so are less likely to be part of what the GenAI tool produces. Keep in mind, though, that the more frequently you repeat an assignment over the years, the more likely it is that past students' responses will work their way into the LLMs that a GenAI platform scrubs.

3. Require specific citations in written work. Sooner rather than later, GenAI platforms will provide this information; but for now, their ability to do this is highly limited. So, asking students to provide detailed and accurate citations will require that they do the work of locating the information they cite.

4. Write assignments about current events, especially local ones.

5. One of the critiques of GenAI's ability to write computer code is that often the output code is illogical or very different from how a human might write that code, even if it functions adequately. This makes it difficult for the code to be incorporated with other code or for a human to edit it. Asking students to edit GenAI-generated code can help them see differences between how humans think about code and how GenAI writes code.

6. Some instructors recommend that you design assignments that center on personal examples and experiences. It's important to keep in mind, though, that many GenAI programs can produce narratives that will at least provide a foundation for a student's answer. For example, a GenAI program might not be able to respond to "write an essay about my experience visiting Brazil," but it will respond to "write an essay about an eighteen-year old's experience visiting Brazil." However, synthesizing personal experiences with specific and narrow classroom objectives can localize assignments such that they would be difficult for a GenAI to provide a complete or cohesive answer.

7. Ask students to engage with GenAI chatbot functions in conversation. Students can learn critical thinking skills this way. GenAIs, particularly those that function with large scale LLMs, are able to return basic responses to questions posed. But, if we teach students to use those basic responses as jumping-off points for follow-up questions, they may come to see the use and value (along with the downsides) of what GenAI provides.

8. Design assignments around fact checking. Have students generate GenAI responses to a prompt and then ask them to work through each claim and statement in the response for accuracy and authentication. Have them provide sources for the GenAI's response. Then, ask students to write about what the bot got right and what it got wrong and the implications of those accuracies and inaccuracies.

9. Focus on *processes* (like the writing process) rather than on *products*.

Our role as instructors must now include teaching students how to engage GenAI tools critically and responsibly. In this way, there are three guiding principles we might take up when teaching students how to use GenAI technologies in their academic work:

▷ **Encourage students to understand that the work they submit must be their own**, not copied in total or in part from a GenAI output. This approach can also help students better see how they might use GenAI to inspire their own thinking. Responses must be their own unless otherwise instructed.

▷ Fundamentally, students should know that if they use a GenAI platform in responding to an assignment, they must be **transparent about that use and document it**.

▷ Students should understand the **recurring flaws** in GenAI content and the risks of using such content, such as inaccuracy, bias, obsolescence, and unclarity of origins.

Your assignment designs should incorporate ways to help students think through why they might turn to GenAI to inform their own writing:

▷ What is your purpose in using GenAI to complete this assignment?

▷ Does using GenAI help you learn the lessons the assignment is designed to teach?

▷ Will using GenAI save you time in responding to the assignment, and, if so, what do you gain or lose in saving that time?

▷ Does using GenAI limit how you would otherwise respond to this assignment?

▷ Is GenAI the right tool for this assignment, or are there other tools better suited to the task, such as a search engine?

▷ Are you using GenAI in ways that are ethical and that do not conflict with your instructor's (or your school's) academic integrity policies?

Generative AI and Writing Instruction

Almost immediately following ChatGPT's public debut in late 2022, questions about the effect of GenAI on writing instruction emerged as a key concern among educators. Two primary questions drove significant parts of the surging conversation. First, if GenAI platforms could produce responses to writing prompts that could pass as being written by humans, how would we police their use and enforce plagiarism policies for GenAI manufactured texts?

The second question is more philosophical or speculative: if GenAI technologies can produce writing that effectively communicates information within specific contexts, does higher education need to continue to emphasize writing instruction to the degree it does? Keep in mind that the first-year writing course—"Freshman English"—has been a near-universal course offering for all incoming college students for more than half a century.

These questions are both interesting and important. They reveal an undeclared yet crucial part of how we talk about GenAI and higher education: **GenAI technologies are now inextricably part of how students will write in the academy and beyond, and as instructors and administrators, we will necessarily have to adjust our pedagogies, policies, learning outcomes, and objectives when teaching writing.**

To that end, there are several ways to think about the connections between GenAI, writing, and writing instruction. The following considerations are meant to initiate conversation among faculty, administrators, and researchers, not to provide definitive or

comprehensive methods or solutions. These are only four concepts with which to begin thinking and talking about writing instruction on your campus; there are many others.

1. *GenAI for Invention:* When instructors began to recognize that banning Wikipedia in the writing classroom was untenable, many began to encourage students to turn to Wikipedia pages as part of their invention strategies. They encouraged them to look for key concepts, to review the citations listed, and to use the information in a wiki page as a way to generate ideas. Can GenAI serve a similar purpose in a writing class? Are there ways to encourage students to use GenAI to generate ideas, to point them toward larger discussions about a subject, to better understand what a writing prompt requires? **Can we find ways to use GenAI to help students develop their own avenues into a conversation?**

2. *GenAI for Revision:* Composition scholars have long professed the idea that writers improve when they receive critique throughout their writing process. This is where familiar practices of peer revision and writing assessment emerge, as both are methods for giving students feedback about their writing. They also teach students how to improve their own writing by giving feedback to others. **Can GenAI provide a similar function, adding a new phase of feedback to student writing processes?** Students already rely on AI technologies like spell checkers and grammar checkers to "fix" their writing before submitting assignments. Could GenAI provide a more encompassing level of feedback for student writers? Such questions necessitate, too, that we ask how integrating such practices into student writing might affect the ways in which students approach writing tasks after they leave college.

3. *GenAI for Critical Thinking:* Composition scholars have long argued that one of the central objectives of teaching writing is to show the inextricable connections between writing, problem solving, and critical thinking. One of the most prevalent critiques of GenAI in education has been the claim that GenAI will generate a new degree of laziness among students; that if GenAI platforms can do the work for them, students won't take the time to think about the materials or assignments they've been given. Of course, as noted earlier, similar critiques have been levied throughout the history of writing. Instead, GenAI is likely to provide a moment in which to rethink what we mean by critical thinking and problem-solving. The current state of GenAI technologies is such that we know that GenAI cannot assess the value or accuracy of the information it provides. Thus, teaching students to be critical of GenAI-generated information must be central to the critical thinking skills we teach. Likewise, we have traditionally taught writing in colleges as a method for teaching students how to demonstrate knowledge. **Working with GenAI technologies can assist students in understanding the differences between their own presentations of accumulated information and GenAI's.** The same concept holds true for teaching students about rhetorical organization of information.

4. *GenAI for Research:* Current iterations of GenAI are incapable of discerning between accurate and inaccurate information. While this provides opportunities to teach students critical skills for evaluating information, it also provides opportunities to teach students critical research skills. Just as we used to teach students how to navigate a print-based library, and now teach students how to be alert to accuracy in web-based research, **we will need to teach students to be critical of GenAI-generated content they intend to use.** However, as GenAI becomes more reliable and accepted (as has Wikipedia), we will also need to teach students how to use GenAI as a tool for conducting reliable research.

As institutions deliberate about GenAI on their campuses, those conversations will benefit from including writing studies experts and writing program administrators at all phases of discussion, from early inquiries to policy development. Given that writing-intensive courses will inevitably be proving grounds for the GenAI debate, having experts in writing studies—particularly those who work in digital writing and related fields—will be invaluable. The same should be said for including humanities scholars (and digital humanities scholars specifically), as institutions begin developing GenAI policies.

▶ **WHILE MOST** conversations about GenAI and writing instruction have thus far focused on students' engagement with GenAI, we need to consider, too, how GenAI technologies stand to affect the work we do as instructors, as well. Can we, for example, use GenAI platforms to complete some of the tasks that writing instructors repeatedly perform? Is there benefit in using GenAI to provide students with fundamental information about writing? For example, rather than taking classroom time to explain the use of a semicolon, can a GenAI platform fulfill that and other similar tasks for us? That is, are there teacherly tasks that GenAI might expedite? Of course, these are the sorts of questions that also drive debates around machine grading of student papers, but they are the exact kinds of conversations we need to have about GenAI.

Product Versus Process (Same as It Ever Was)

In the 1960s and 70s, writing instructors began focusing on the process by which writers produce written text. The now commonplace idea of the writing process (prewriting/invention, research, drafting, revising, and editing/proofreading) gained popularity because rather than reducing writing instruction to correcting student writing—that is, showing students what they did incorrectly in their writing—instructors could teach students strategies for working toward a product. The "writing process" approach became prevalent as a pedagogical method because it is teachable, self-reflexive, recurrent, and replicable.

GenAI does not so much disrupt the process model for teaching writing as it simply ignores it. As explained earlier, GenAI receives a prompt, scrubs its LLM, and produces

an output. GenAI cannot, however, explain how it got from the prompt request to the product. It doesn't prewrite, produce drafts, or revise those drafts. Nor can a GenAI produce multiple drafts of an essay or other document, improved with each iteration (barring additional input). Because higher education emphasizes students' understanding of how they move from inquiry to product, the black box logics of GenAIs confound educators. There is no way to assess how a GenAI gets from point A to point B because GenAI cannot describe the route between the two. Thus, conversations about GenAI and assessment of student work have rejuvenated focus on processes. Students in Western educational systems are browbeaten by the mantra "show your work" because our educational philosophies espouse the idea that providing an answer—a product—is only part of learning. Being able to explain how one got the answer reflects a deeper learning, and, possibly, greater retention and transferability.

That being the case, it becomes more valuable in the age of GenAI to refocus assessment on *processes* rather than *products*. Asking students to be self-reflexive and to show how they moved through a writing process, or the steps they took to solve a mathematical equation, redirects the learning from product to process. In this way, the product is of less value within assessment rubrics than is how the student ended up at the product. Such strategies also emphasize students' learning and their ability to explain why and how they made the choices they did.

03

Artificial Intelligence and Administrative Policies

Policy and Prohibition

Less than two months after ChatGPT's launch, New York City's Department of Education banned it from use on school networks and devices. The Los Angeles Unified School District soon followed suit. Similar responses unfolded internationally, including in school systems across Australia and Italy. Colleges and universities have been less eager to initiate institution-wide prohibitions. However, as of March 2023, 31 percent of college students surveyed claimed that their instructors, course materials, or school honor codes explicitly prohibit using AI tools. Realistically, given the rapid evolution of GenAI technologies and their ever-growing use in many settings, such bans are likely to be short-lived.

There are two fundamental principles that inform knee-jerk bans on GenAI technologies in educational institutions. First is the idea that students will inevitably use GenAI to cheat. This position places policing at the fore of education, above teaching and learning. It also casts students as inherently nefarious, rather than curious and inspired. No prohibition on using GenAI technologies will stop students who decide to cheat from cheating. Second, and perhaps more important, **banning student engagement with GenAI tech will leave students and instructors ill-prepared to work with GenAI technologies beyond the classroom in their professional, civic, and personal lives.** While bans might seem a quick fix approach, prohibitions and blanket policies do more than restrict students from using GenAI technologies to cheat: they establish positions that restrict the possibilities of education and research.

▶ **DESPITE OUR DESIRE** to talk about our students as digital natives (a deeply problematic term itself), we forget that their digital experiences are rarely grounded in academic, professional, civic, or even personal experiences or expertise. That is, students need to learn to use the platforms, programs, and applications needed to complete tasks just as they need to learn any other tool or concept. When it comes to GenAI, the situation is the same: they have not yet learned to use these tools in useful, effective, and ethical ways. That is, digital natives are not inherently native GenAI users any more than they are native users of CAD, Photoshop, or any other application. It is our task as instructors, researchers, and administrators to assist them in learning how to use these tools responsibly.

Certainly, academic integrity policies need to be revised to account for GenAI technologies. However, revisions must provide degrees of flexibility that define how using GenAI can violate academic integrity policies like those about plagiarism or collusion, while acknowledging that such tools can be used responsibly.

Remember, too, that prohibition always leads to some degree of clandestine use, rather than complete adherence. That is, whether banned or not, some students will still use the tools; they already do. While unsanctioned uses of all kinds of technologies have led to remarkable discoveries and innovations over the years, sanctioned and monitored uses have led to more. Prohibiting the use of a tool is less productive than designating when and how students should use it and then teaching them (and ourselves) the ethical criteria for such use.

Given the wide variety of learning objectives and teaching methods from one discipline and instructor to the next, it seems most productive to design policies about GenAI that allow local contexts to dictate how GenAI is or is not permitted in a given class. That is, instructors who have philosophical, methodological, or pedagogical reasons for not wanting students to use GenAI in their classes should be provided with the authority to include statements in their syllabi prohibiting GenAI. However, instructors who want their students to engage with GenAI should also be given the autonomy and authority to integrate GenAI into their course and assignment designs.

As administrators, researchers, and instructors, our approaches to developing new policies about GenAI need to address institutional-level considerations of academic integrity, and they should allow individual instructors to determine if and how GenAI can be used in specific classrooms. Removing that kind of autonomy from a given instructor restricts them from teaching to their strengths.

The fundamental policy that we should adopt isn't one holding GenAI to be universally prohibited or universally permitted, but rather one holding that all use of GenAI should be transparent and documented.

It may be a while before GenAI technology reaches a plateau. Until that time, it's unlikely that we'll be able to develop consistent, formal standards for AI use in education. Thus, for now, the best resources we have for thinking about how and why we might develop certain policies will fall to the instructors and students who engage with GenAI in the contexts of specific curricula, with their own disciplinary customs and local objectives. This is why the continued conversations about the values and risks of GenAI in higher education must be encouraged rather than stifled by blanket prohibitions. Innovation will come from engagement.

Career Readiness

Career readiness has become central to the missions (and expectations) of contemporary higher education institutions. From career resource centers to pro-seminars and professionalization programs, college education has embraced the demand to prepare students for the workplace as part of its core mission. This is particularly true for vocational, technical, and trade schools. As more and more industries embrace GenAI as part of their daily operational procedures—from CEOs using GenAI to generate reports or answer emails, to graphic designers using GenAI to create product images for catalogs or to develop technical schematics—higher education must now consider how GenAI will affect the ways in which institutions prepare students for the workplace. **GenAI skill is now inseparable from career readiness and workplace readiness, and will become central to our efforts in higher education.**

Career readiness programs prepare students with the skills needed to locate, acquire, retain, and grow in their careers. As such, career readiness is central to nearly every major offered on any campus. Understanding how GenAI interacts with specific majors, and with the careers students in those majors might pursue, should now be understood as a core component of career readiness in higher education.

Career readiness generally focuses on helping students acquire three key types of skills: academic skills, employment-ready skills, and technical skills. Academic skills include the ability to communicate clearly and effectively (particularly in written media), as well as basic research skills, and the ability to perform basic computations. Employment-ready skills are those seen as necessary for workplace readiness; they include critical thinking skills, problem-solving skills, collaboration and teamwork skills, creativity, responsibility, and integrity. Technical skills are the technical abilities relevant to job-specific contexts and vary greatly by discipline and career.

GenAI technologies, as we have seen throughout this resource, now play a significantly larger role in each of these three contexts of skill. Any consideration about the integration of GenAI into higher education should include significant consideration of your campus's career readiness objectives and resources. Representatives from career readiness programs should be included in any discussion of institutional policies regarding GenAI.

AI and Your Institution's Mission and Strategic Goals

Most institutions of higher education have similar missions and strategic goals. These objectives are usually reduced to research, teaching, and service, but higher education's missions are much more complex than that. These objectives are inextricably connected with those of solving societal problems, preparing workforces, and anticipating what comes next. Given the rapid emergence of GenAI within our society, and its bearing on the workforce, each institution should reflect on how to approach GenAI within the parameters of its mission and its strategic objectives. All post-secondary institutions need to establish methods for teaching students how to responsibly interact with GenAI technologies, develop conceptual and applied uses for GenAI, and share their knowledge with local and global communities.

It is the responsibility of institutions to provide the resources that faculty, administrators, and students need to develop responsible and rigorous ways to engage GenAI within their institution's teaching, research, and service missions. Such resources should begin with faculty training, and with research opportunities to work with GenAI technologies. Given the anxieties that the release of ChatGPT has triggered among many educators, it seems judicious to formally prepare instructors—no matter their rank, teaching obligations, or contractual assignment—to engage their students on issues relating to GenAI. Such preparation might be treated as continuing education or professional development. As such, faculty should be compensated for this professional development in support of their institution's mission.

Treating GenAI integration as intrinsic to your institution's mission can also facilitate collaboration among campus leaders. This can help promote centralized relationships as well as collaborations among previously unconnected entities on campus. Since GenAI integration is not discipline-specific but rather a cross-campus matter, institutional infrastructure must account for whole-campus integration.

Wherever strategic goals are being adopted to better address the role of GenAI, it's essential to include all entities that serve large swaths of the student body, such as libraries, honors programs, and writing centers.

AI and Your Institution's Digital Literacy, Competency, and Fluency Efforts

Just as GenAI engagement should be situated within an institution's mission and strategic objectives, so too should it align with institutional digital literacy, competency, and fluency efforts. More and more institutions are developing long-term goals that include the advancement of digital literacy skills and requirements, and AI literacy needs to be a part of this. Given the rapidly expanding influence of AI on our daily lives, students should learn basic AI functions as part of their education. Expected degrees of fluency, competency, and literacy might best be determined both at the institutional level and

within the discipline-specific settings of individual majors. And it should be noted that what qualifies as AI literacy need not—at least in some disciplines—be highly technical. The authors of a recent article on this topic define AI literacy as "the ability to understand, use, monitor, and critically reflect on AI applications without necessarily being able to develop AI models themselves."

Interestingly, **many of the efforts to bring AI literacy to higher education have not emerged from computer science, IT, or related disciplines, but instead from areas such as medicine, business administration, education, and the humanities.** This is, in part, due to the fact that these disciplines—among others—have experienced the most rapid increases in the applications of AI to their academic and workplace settings. As such, there is both a pedagogical imperative and an ethical obligation to prepare students in understanding how AI functions in those specific contexts. Again, locality is important here: allowing individual programs and departments to adapt AI literacy approaches to their specific disciplinary contexts will provide students with AI literacy skills that are best suited to their academic and professional goals.

Accessibility, Diversity, and Bias

As higher education struggles to understand the risks and possibilities associated with GenAI, questions regarding Diversity, Equity, and Inclusion (DEI) inevitably surface as central. Like other emerging and evolving technologies, GenAI has been cast as posing a risk to DEI efforts in one breath and as a powerful tool for DEI in another.

The crucial observation is that we cannot address GenAI in higher education without significant consideration of DEI. As with many other issues, most conversations about GenAI and DEI must necessarily be local. However, several universal factors should be brought to the fore in any such conversation:

▷ *Access.* As with many other digital technologies, access to GenAI is not equally distributed, nor is it likely to be so any time soon. Not all students have access to GenAI technologies. These gaps can be the result of socio-economic or geographic factors that limit access to broadband Internet and compatible devices, or they can stem from a lack of background or training in the use of such devices. While ChatGPT doesn't currently charge for access to its most basic service, many of the other GenAI platforms do, and some are releasing with subscription-only access. We cannot assume that all students will have equal access to GenAI at present or going forward. Access is inextricably bound to affordability.

▷ *Exclusion.* Because a GenAI platform can only deliver output based on the information it accesses through its training and its LLM, significant information is always excluded. In particular, large quantities of data created by marginalized people are often excluded from digital databanks. The narratives of many disenfranchised groups are not readily accessible to LLMs because their voices and

narratives have been historically excluded from our accounts of history. As such, their contributions frequently remain absent from LLMs and, thus, from what GenAI engines produce.

▷ *Bias.* Just as GenAI is unable to assess value or accuracy, it's also unable to identify bias in the information it scrubs and re-presents. As such, the algorithms that drive a GenAI lead to biases that it employs in its search for and repetition of information. For example, when ChatGPT first appeared, many users found that it would readily provide left-leaning responses to prompts; but when asked to convey information from right-leaning political positions, it stated that it was unable to do so. Likewise, several GenAI platforms have been shown to draw their data from sources with some political affiliations but not from those with opposed affiliations, thus returning responses that reflect this bias. Even if a GenAI could draw data from a variety of political stances, it wouldn't be able to appropriately weigh these competing positions because it can't assess the value of each. In short, GenAI technologies are not apolitical, as their algorithms and LLMs tend to be biased toward a particular perspective.

▷ *Language diversity*. Keep in mind that—for now—many GenAI platforms only operate in English. While efforts are underway to develop multi-lingual and other-than-English platforms, GenAI is primarily an English-speaking technology. As such, GenAI currently works in homogenous linguistic models, reinforcing the patterns, styles, and word choices of a single language. Such linguistic homogenization, we know, restricts diversity, as the debates regarding Standard Written English have exposed.

Note, too, that concern over language diversity extends beyond the written output of ChatGPT (which, notably, is growing in its ability to read and access content in languages other than English) to questions of image output, as well. One study, for example, reports that when Dall-E2 (a popular GenAI image generator) was asked to create an image of a "close up photo of hands typing on a laptop," the platform returned four images of masculine, Caucasian hands. Only when prompted to create an image of a person of color's hands did it do so. If GenAI platforms default to such patterns, then diverse perspectives are muted.

These—and many other—matters of diversity, equality, inclusion, access, and bias must be central to how we consider the role of GenAI in higher education.

Other Concerns

There are three other factors to which administrators, instructors, and students should be alert:

1. **Privacy:** GenAI learns from each interaction to which it responds. So, part of the function of such platforms is to gather more data so as to improve functionality.

Some platforms, including ChatGPT, also gather data about the user. For example, ChatGPT's privacy policy states that it gathers multiple types of data:

▷ Account Information: When you create a ChatGPT account, you must agree that its creator OpenAI can store "your name, contact information, account credentials, payment card information, and transaction history."

▷ User Content: When you use ChatGPT, you allow it to collect "personal information that is included in the input, file uploads, or feedback."

▷ Communication Information: In communication with OpenAI, you agree that the organization "may collect your name, contact information, and the contents of any messages you send."

▷ Social Media Information: If you interact with OpenAI's social media pages, you allow it to collect "personal information that you elect to provide…, such as your contact details." OpenAI also receives aggregate reports about usage from those social media sites.

▷ User Data: ChatGPT retains information provided by your computer, including log data (such as your Internet Protocol address, browser type, and time of usage), usage data (such as "the types of content that you view or engage with, the features you use and the actions you take"), device information, cookies, and analytics.

2. **Intellectual Property:** When a GenAI responds to a prompt, the information it provides has been gathered from the data to which the GenAI has access. While many GenAI developers are working to have their platforms identify and cite the resources from which they gather data, most currently do not. Thus, GenAI applications may provide you with information that they do not have permission to distribute. As such, the platforms themselves may violate intellectual property laws.

Additionally, because most GenAIs retain information input by their users, they may gather, store, and recirculate this content to other users. This content may include the user's intellectual property, as well as any personal information they input to the program.

3. **Longevity:** One of the functions of most GenAI programs is to expand their LLMs so as to generate better responses in future. Since GenAIs retain all inputs and outputs, users should be cautious with their prompts. It may be funny now to ask a GenAI an offensive or lewd question, but one should consider the ramifications if that same platform later integrates this input when generating a response to someone else's query.

04
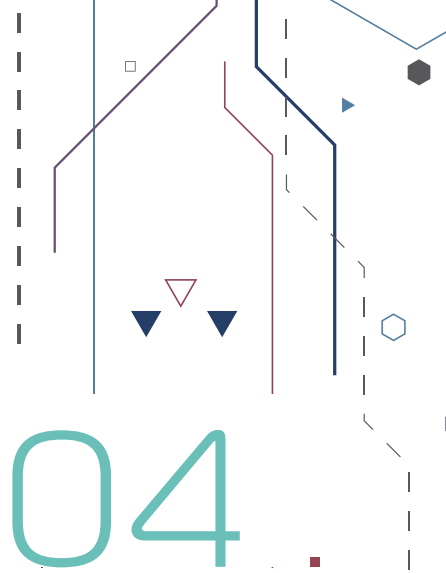
Best Practices

Strategies for Developing Best Practices

Rather than provide specific best practices, I offer some considerations upon which to develop local best practices and start productive conversations on your campus.

▷ *Stakeholders:* Because GenAI technologies are gaining traction in just about every discipline, **it is important to include representation from across academic disciplines when developing best practices.** When it comes to campus technologies, many are tempted to turn to IT and computer science experts alone. However, given the broad influence of GenAI across disciplines, it's critical to seek input from as many divisions as possible. To that end, involving representatives who work specifically in writing, medicine, business, and other areas where GenAI technologies are establishing footholds will be beneficial. Libraries, honors programs, student success programs, career preparation centers, and the like should also be represented in these discussions, as should administrators (IT administrators in particular). Students should also be involved, especially those who've engaged with GenAI in their studies and research.

▷ *Scope:* **Determine which aspects of AI and GenAI you want to address in your best practices.** Will your best practices cover student use? Faculty use? Administrative use? Research use? Will your best practices address GenAI only, or broader aspects of AI in general? (Given the massive range of AI, and the assumed urgency of GenAI, it may be best to separate these topics in discussions of best practices.)

▷ *Governance:* Think about how best practices might be enforced, and how this would translate into local and institutional policies. **Consider establishing a GenAI governing board made up of the stakeholders described above.** There are different ways that such a structure could work in relation to curriculum development, research policies, or technical policies. When applicable, involve student government and student honors organizations as well.

▷ *Use:* Best practice policies will only work if you have some idea as to how GenAI is *already* being used on your campus. **Survey faculty, students, and administrators about how they use GenAI in their work.** Those who already use GenAI may have insight into how it will be used in the future. Include, too, individuals whose research already examines GenAI, and instructors who are vocally in favor or vocally opposed.

▷ *Concerns:* **Best practices need to respond to the concerns of faculty, students, and administrators.** Surveying stakeholders can help identify the concerns your campus community members have. You should then address those concerns—or at least the most prevalent ones—when developing best practices. Such a survey might reveal areas in which policies are most urgently needed, such as those concerning plagiarism and how to address this in the classroom.

▷ *Legality:* Inevitably, GenAI use is going to expose a range of legal issues, including issues of data privacy, intellectual property, cyber security, bias, discrimination, validity, misinformation, and so on. **Consider the various legal issues, and their relative priority, as you develop best practices.** Some legal issues can be anticipated and aligned with existing policies; it may be beneficial to look for these overlaps and connections.

▷ *Flexibility:* Because GenAI is relatively new and still evolving at a rapid pace, **best practices should be developed generically and flexibly.** As these technologies and the norms surrounding them fall into more consistent patterns, best practices can be honed. Thus, policies should be organic and in perpetual development and revision.

Conclusion: A Parable of the Luddites

During the early nineteenth century, the British textile industry rapidly mechanized, turning to machines to produce textiles previously manufactured by hand. The Luddites were a secret organization of textile workers who opposed such automation and began destroying textile machinery in protest. They claimed that the new textile machines allowed manufacturers to avoid standard labor practices and that those who used such machines did so in "a fraudulent and deceitful manner." They feared, too, that textile machines would leave traditional workers without jobs. Luddites worried that their skills and the art of their crafts would be lost as textile machines became

the new standard. From 1811 through 1816, the Luddites led a region-wide rebellion against the mechanization of their industry.

As we all know, mechanization won out in the textile industry, and most consumer textiles are now manufactured mechanically. However, the widespread mechanization of the textile industry did not eliminate the traditional textile worker. Instead, some of those workers found jobs in the automated factories, lending their expertise to the new mechanical processes; some even contributed to the evolution of mechanized production. Many of the traditional textile manufacturers continued their craft and passed on their trade knowledge to successors. The craft of textile manufacturing did not vanish entirely as automation became the norm.

Much like the Luddites, many will resist GenAI in education. Many will wish to ban these technologies; many will claim that they displace human thought and skill. (In fact, as I was working on this resource, I received an email from an alumnus of our program, asking what our department is doing "to protect writers from artificial intelligence".) Yet, what many do not recognize is that the mechanization of the textile industry did not render the skills and practices of the textile workers obsolete. Instead, it created opportunities for some craftspeople to take their expertise and experiences into the mechanized age. It also opened spaces for artisan textile manufacture, which established variations in quality expectations, craftsmanship, and prestige. Despite their resistance, many Luddites found paths to success within the textile industry by adjusting their understanding of the role of their work.

The story of the Luddites makes evident that there are possibilities for adapting our expertise to new technologies. If we apply our expertise to these technologies, even if only minimally, we can encourage them to develop in ways that are useful to us, rather than at odds with our interests. It's also important that we take into account resistant voices and not dismiss or ostracize those faculty who may not be willing to embrace GenAI for any number of reasons. The role of GenAI in higher education is still being determined, and though it's not possible for higher education to block its growing influence in the world, we can and should involve a wide variety of perspectives and stakeholders as we forge a path forward.

AI and Writing textbook
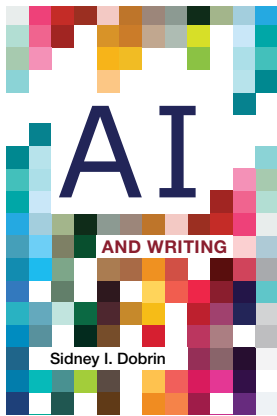
Scheduled for publication in Summer 2023 and available for Fall 2023 courses, *AI and Writing* provides a road map for students and instructors who wish to talk about GenAI programs such as ChatGPT but who don't know where to begin. Writing prompts, discussion questions, and provocative short readings guide students in their consideration and use of GenAI. The focus is on the pragmatic: Can these programs be used ethically in writing? What tasks do they work well for? Where do they fall short? Exercises covering the most commonly taught modes of composition make this book an ideal choice for any instructor who wishes to integrate coverage of GenAI into their undergraduate writing classroom.

Click here to request an advance examination copy or for further information.

Key Features:

▷ Ideal introduction to AI for any writing-intensive course
▷ Writing and discussion prompts address practical and ethical questions of AI use
▷ Provides expert guidance for instructors approaching these topics for the first time
▷ Includes numerous exercises and writing prompts

Contents:

I. **Introduction: ChatGPT and the Generative Artificial Intelligence Surge**
    Automated Writing: It's Not Really New
    History of Writing Technologies and Cultural Panic

II. **What Is Generative AI? (and for that matter, what is AI?)**
    The Myths of Artificial Intelligence
    Artificial Intelligence Briefly Defined

Generative Artificial Intelligence Briefly Defined
How Does Generative Artificial Intelligence Work?
What are our Assumptions and Expectations about GenAI?

III. Generative AI and Academic Integrity (as well as professional, civic, and personal integrity)
Generative AI and Integrity
The Plagiarism Problem
Spaces of Judgment

IV. What Is an Author?
Generative AI and the Idea of Authorship
Where Is My Writing?
Writing with Algorithms
Does Generative AI "Write"?

V. The Places of Generative AI Writing
Academic
Professional
Civic
Personal

VI. Generative AI and Writing Processes
Generative AI and Invention
Generative AI and Drafting
Generative AI and Revision

VII. Generative AI and Creativity
What Is Art?
What Is original?
Visual Rhetoric and GenAI
Generative AI and Multimodal Writing

VIII. Generative AI Best Practices

About the Author

SIDNEY I. DOBRIN is Professor and Chair of the Department of English at the University of Florida, and for ten years directed its writing program. He is the Founding Director of the Trace Innovation Initiative, serves as a Digital Thought Leader for Adobe, and is a member of the Florida Institute for National Security, which is part of the University of Florida's Artificial Intelligence Initiative. Dobrin is the author and editor of numerous books and articles about writing, technology, and ecology.

Sidney I. Dobrin is available to consult with post-secondary institutions about Generative AI, including policy development, curricula, and the development of faculty training programs. Please contact him for further information.

Glossary

Algorithm
A process or set of rules or instructions used for solving a problem or calculation.

Applied AI
How to use AI, both generically and specifically.

Artificial Intelligence (AI)
The theory and development of computer systems that can perform tasks that previously required human intelligence.

Conceptual AI
How we think about the ramifications of developing AI.

Deep Learning
A series of complex algorithms that are modeled on the human brain and the structures humans use to think through tasks.

Generative Adversarial Network
A variation on a machine learning algorithm that allows AI neural networks to take the data gathered and learned from machine learning and use that data to generate new iterations of that information through written text, image, and sound.

Generative Artificial Intelligence (GenAI)
A set of algorithms that can generate seemingly new, realistic content, including text, images, or audio, based on training data and Large Language Models. GenAI can be thought of as any artificial intelligence that can produce what appears to be original content.

Large Language Models (LLMs)
Machine learning algorithms that use deep learning to identify, scrub, and process natural language. LLMs "train" using large quantities of text data to learn patterns and entity relationships in the language.

Machine Learning
The ways in which computer systems use algorithms to analyze and draw inferences from patterns they identify within specific data sets.

Neural Network
A mathematical system that scans large amounts of data to identify patterns.