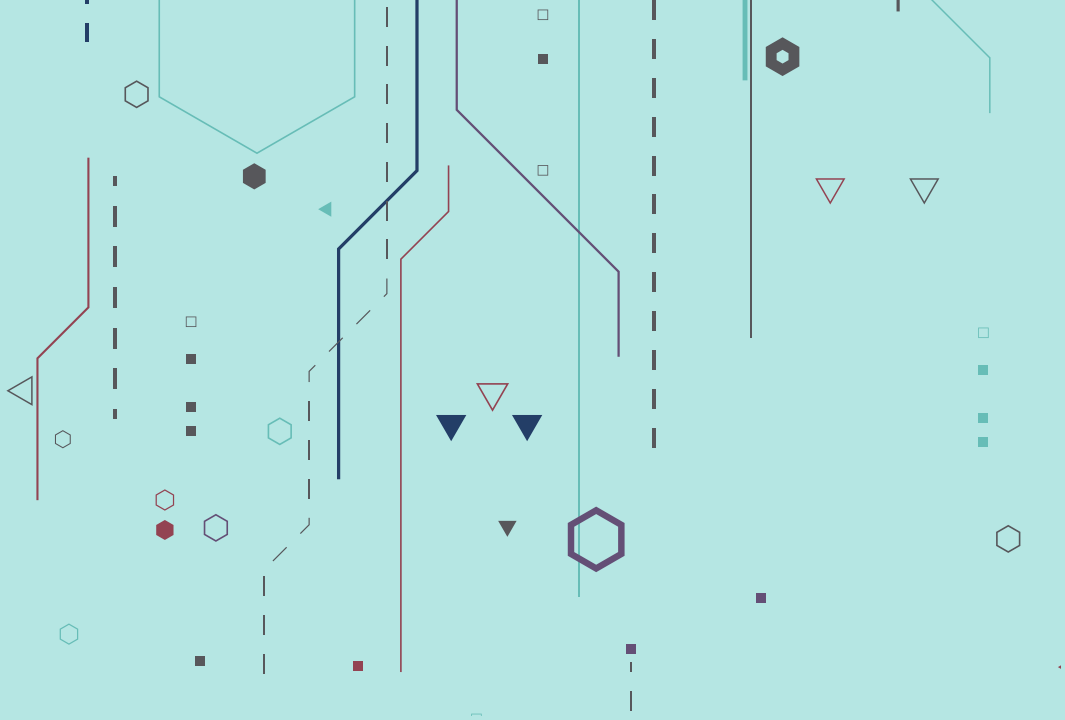

Talking about Generative AI

A GUIDE FOR EDUCATORS

Sidney I. Dobrin

Version 2.0

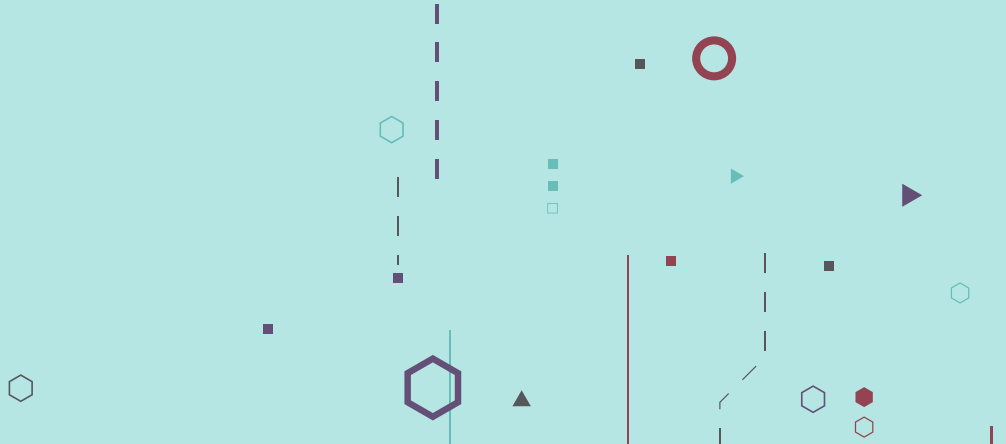

Contents

Goals of This Resource 4

01 What Is Generative Artificial Intelligence? 5
- How Does Generative AI Work? 7
- Generative AI and Higher Education 9
- History of Writing Technologies and Resistance to Them 10
- Generative AI Across Disciplines 12

02 Artificial Intelligence and Administrative Policies 15
- Policy and Prohibition 15
- AI and Your Institution's Mission and Strategic Goals 17
- AI and Your Institution's Digital Literacy, Competency, and Fluency Efforts 18
- Accessibility, Diversity, and Bias 19
- Privacy 20

03 Artificial Intelligence for Course Instructors 23
- Curricular Relevance (or, Why Integrate GenAI?) 23
- Generative AI and Integrity 25
- Authenticity and the Plagiarism Problem 26
- Designing Assignments that Outpace GenAI? 29
- Assignment Design in the Here and Now 30
- Generative AI and Writing Instruction 32
- Generative AI and Reading 35
- Generative AI and Instructor Labor 36

04 Artificial Intelligence and Assessment 39
- Why We Assess 39
- Reimagining How We Assess 41

05 Artificial Intelligence and Student Career Readiness 44
- Workplace AI Literacy 45
- Learning to Learn 46
- Industry to Curricula 47

06 Faculty Development 50
Faculty AI Readiness 50
Faculty Needs 51
Professional Development Programs 52
AI Literacy for Instructors 53
Labor and Job Security 54

07 Infrastructure 57
Technological Support 57
Training and Development 59
Compliance 60
Sustainability 61

08 Best Practices 63
Strategies for Developing Best Practices 63

Conclusion: A Parable of the Luddites 65

***AI and Writing textbook* 67**

About the Author 69

Glossary 70

Goals of This Resource

The introduction of ChatGPT in November 2022 served as a wake-up call for higher education, and as often happens when an alarm sounds, some people are startled and frightened, some are angered, some are slow to react, some ignore it altogether, and some embrace it. Generative AI (GenAI) 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 in addition to exploring how AI will affect the work of higher education.

Talking about Generative AI 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.

That said, this resource is not designed to recommend specific AI platforms for adoption. As more and more platforms are developed for specific uses, individual institutions will need to make choices about what applications to provide on their campuses and what guidelines will direct their adoption policies. This resource does attend to important considerations relating to adoption, and while I occasionally refer to specific AI platforms throughout, I do so as a matter of example; my naming of a particular platform or company should not be read as an endorsement.

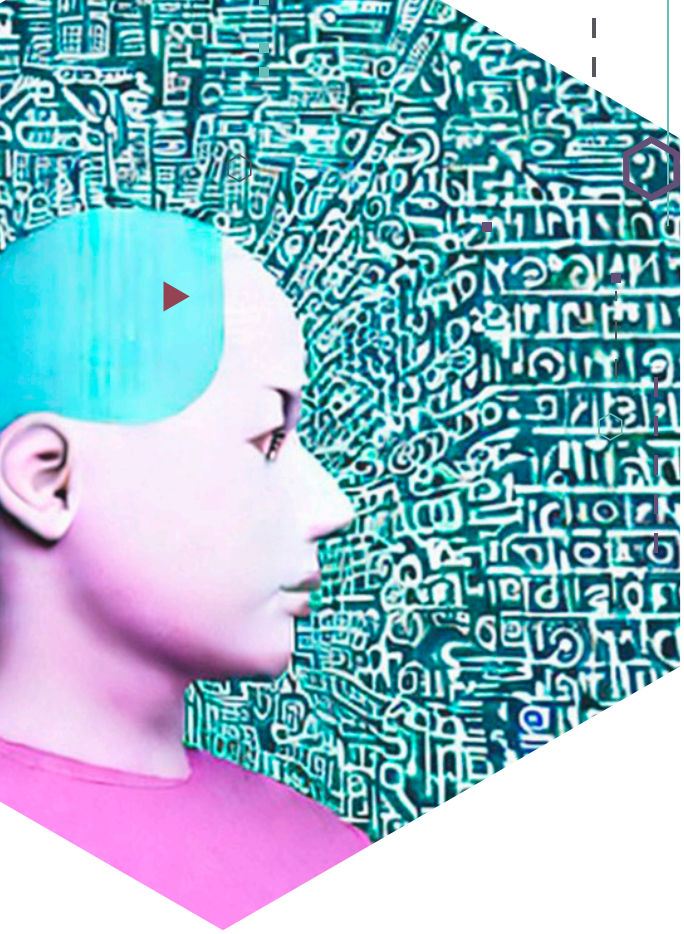

01

What Is Generative Artificial Intelligence?

In this resource, I'm not going to provide a comprehensive history of artificial intelligence—there are several great resources out there that do this: see, for example, Michael Wooldridge's *A Brief History of Artificial Intelligence* (2021)—nor am I going to unpack all the nuances of the various functions of AI and GenAI. However, it is 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

THE LANGUAGE we apply to understanding how AI works tends to be anthropomorphic: *learning*, *thinking*, *training*, etc. While this language can help us conceptualize complex technologies, attributing human characteristics and behaviors to AI gives rise to philosophical questions about AI algorithms' actual ability to "learn" and "think" like humans. We need to recognize two key parts about AI's potential: First, despite the doomsday narratives that popular culture posits (Skynet, Hal, Cortana, WOPR, and so on), the truth is that the hardware necessary for AI to achieve such levels of self-awareness doesn't exist at this time. The materiality of the AI machine limits what capabilities an AI can actually have. Second, we have to recognize that although AI's "thinking" may be effective in specific contexts, it is not the same as human thinking. It is a different—often not understood—algorithmic process. Recognizing our human tendency to anthropomorphize, along with being able to identify the differences between conceptual and applicable approaches to AI, can help us develop our understanding of AI use in higher education.

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 about *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.

One 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 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 datasets. That is, when an AI recognizes patterns within a dataset, it "learns" to make inferences about those patterns. Machine learning requires AI systems to "train" using predetermined datasets 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 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.

How Does Generative AI Work?

Researchers made only minimal advances in developing AI between the early explorations of the 1950s and the early 2000s. Then, in 2009, advances in deep neural network development rapidly enhanced speech and vocabulary recognition. 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 whales, 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 the **generative adversarial network** (GAN), 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 datasets, 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 datasets 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.

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 dataset. The algorithm has "learned" to identify specific characteristics within that pattern.

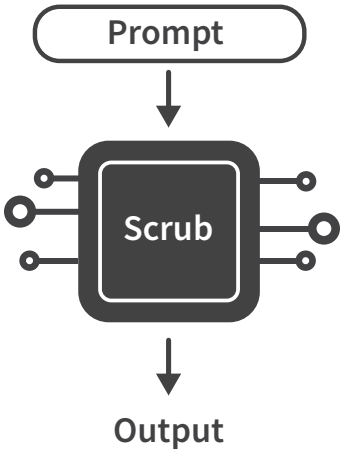

GenAI works like this: a user provides the AI with a prompt asking the AI to create a specific deliverable—an essay, a song, an image, the solution to a math problem, a computer program, or the like. The AI then scrubs through all of the data available, looking for patterns and recurring information 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). Platforms like ChatGPT, Copilot, and Gemini are just more complex chatbots, pretrained to locate data, transform that data, and generate new ways of conveying it.

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, the evolving understanding of what a campus library provides, and even the notion of proprietary access to data. 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 Chapter 2 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, 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; these skills remain crucial to teaching responsible use of technology, problem solving, critical thinking, and creativity.**

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 that is widely and digitally available.

Generative AI and Higher 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 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. The platform's release—and the subsequent availability of platforms like Gemini and Copilot—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 not, and

TIME TO 1 MILLION USERS

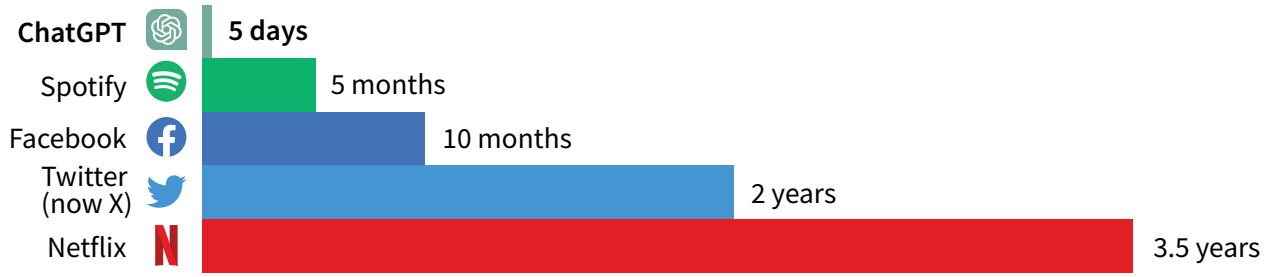

Source: Buchholz, Katharina, and Felix Richter. "Infographic: ChatGPT Sprints to One Million Users." *Statista Infographics*, 24 Jan. 2023.

TIME TO 100 MILLION USERS

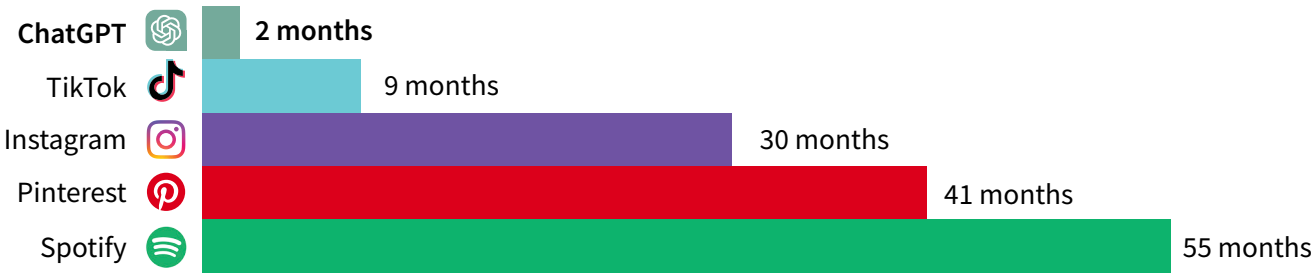

Source: "CHATGPT's on Track to Surpass 100 Million Monthly Users Faster than TikTok or Instagram: UBS." *Yahoo! News,* Yahoo!

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 the Digital Education Council** in 2024, 58% of university students feel that they do not have sufficient AI knowledge and skills, and 72% agree that universities should provide training for students on the effective use of AI tools. Tellingly, the same survey reveals that 80% of the surveyed students report that their university's integration of AI does not fully meet their expectations.

History of Writing Technologies and Resistance to Them

Imagine, if you would, a writing class that bans word processors, spellcheckers, 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. As mathematical statistician Nassim Nicholas Taleb puts it, "Technology is at its best when it is invisible." 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 postliterate 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. That is, GenAI requires that we question our legacy epistemologies. GenAI may seem risky to us because it makes us rethink the things in which we are comfortably entrenched and invested; 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, or what educational research has taught us. 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 others. One of the biggest risks we take in adapting GenAI technologies to higher education is to simply ask how these new technologies can assist us in doing the things we have always done more efficiently. Instead, we now must ask how these emerging technologies can help us do things we hadn't imagined before. We must continuously ask, "What comes next?"

ALTHOUGH MANY of us had been unaware of the use of GenAI bot writers until the recent surge in media attention, AI writers have been churning out content for at least a decade in places we might not expect. The *Los Angeles Times*' Quakebot has been writing news articles about earthquakes since the early 2010s—often much faster than any human could—using data from the US Geological Survey's monitoring equipment. 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: many assume that humans do all the research and writing work that goes into producing it. As in higher education, the notion of integrity is central in the domain of journalism, but this hasn't prevented the latter from assimilating AI technologies; instead, journalism's use of AI technologies has contributed to the very understanding of what it means to conduct research and produce writing. These are the questions that the higher education community, across all disciplines, must now engage with.

Generative AI Across Disciplines

Certainly, college writing has come under scrutiny in the era of GenAI, and writing instructors need to continue to rethink not just how writing is taught and assessed but the very understanding of what writing *is*. However, while popular media has made writing the proving ground for GenAI on college campuses, the truth is that GenAI is being deployed in every academic discipline we teach. Though certainly not a comprehensive list, here are a few examples of ways GenAI is permeating various disciplines in higher education:

- *Art and design:* GenAI platforms provide students with dynamic new ways to think about their work, to approach their creative objectives, and to engage in creative exploration. GenAI is now widely used in many industries for content creation, and teaching students how to use GenAI, with the aim of promoting workplace readiness in creative fields, is becoming more commonplace.

- *Biology, chemistry, and physics:* GenAI algorithms are increasingly being deployed to analyze large scientific datasets because they are remarkably efficient at identifying patterns. They are also valuable for generating data visualizations and textual explanations of those datasets, contributing to how scientists understand and communicate complex data. AI and GenAI systems can also create dynamic simulation scenarios, allowing students to experiment in difficult-to-access contexts in safe, controlled ways.

- *Business, finance, and economics:* AI and GenAI platforms are already becoming mainstream in financial industries. Since business, finance, and economics students will likely use these tools in the workplace, AI business literacies are

becoming necessary parts of the academic curriculum. Students are learning to use AI tools for optimizing investment portfolios and predicting market trends, consumer behaviors, and economic indicators.

- *Computer science:* GenAI tools are used for code generation and debugging in many coding languages; they can provide code structures, offer debugging suggestions, and suggest optimization strategies. Computer science students are now learning not just how to code but how to code in collaboration with GenAI platforms, expediting coding processes and freeing students to focus on higher-order issues.

- *Engineering:* GenAI platforms can provide engineering students with strategies for design optimization, help them understand methods to address things like strength, weight, and efficiency, and allow students to experiment with various materials in simulated scenarios. AI can also be integrated with sensors used to predict potential design flaws, and can be used to develop predictive maintenance strategies.

- *Humanities:* Humanities researchers have begun to use analytic AI tools to conduct evaluations like computational stylistics and sentiment analysis to expand approaches to textual interpretation. Likewise, AI is being used to map ecologies of specific concepts across large bodies of texts, allowing researchers to identify previously unrecognized patterns in historical and cultural data. Some are using machine learning to analyze stylistic details of texts to determine authorship and origin. Others are using AI tools to transcribe and translate texts, making them more accessible to wider audiences. Importantly, AI facilitates interdisciplinary research, bridging gaps between humanities and data science.

- *Law:* AI can assist in legal research by analyzing legal documents, case law, and statutes. GenAI platforms are also valuable for drafting and reviewing legal documents, ensuring consistency and accuracy. As the legal profession expands its use of AI technologies, law students will need to learn appropriate methodologies for engaging with AI.

- *Mathematics:* AI systems can provide math students with personalized practice problems and immediate, step-by-step feedback that adapts to their individual levels of understanding. AI platforms can also be used to explore mathematical conjectures and visualize complex mathematical concepts.

- *Medicine, allied health, nursing, and pharmacy:* In the medical professions AI deployment has become more prevalent than it has in just about any other area of research and practice (fintech is one possible exception). AI is being used for image analysis in radiology, for drug development in pharmacology, and for the creation of personalized patient treatment plans. Students working in these areas will need to learn how to use AI in research and application.

- *Social sciences:* GenAI is being used to simulate complex social phenomena and social interactions (including reflexive approaches to studying the dissemination of information). Using AI in these ways can help students and researchers better understand social dynamics and test theories in virtual spaces. Students in the social sciences can also benefit from learning to use AI algorithms to analyze large datasets, particularly when analyzing unstructured data.

GenAI is increasingly being used across disciplines, and students require appropriate instruction to understand how and why to use these tools both in their schoolwork and in their future workplaces. As more platforms designed for discipline-specific work will continue to emerge, administrators, researchers, and instructors will need to address policies for approving which platforms might be deployed on campus and in specific classrooms. It is likely, in my opinion that platform adoption will soon become similar to textbook adoption, and colleges and universities will need to develop platform adoption policies akin to textbook adoption policies.

02

Artificial Intelligence and Administrative Policies

Policy and Prohibition

Soon after the release of ChatGPT in November 2022, many colleges, universities, and school systems reacted by implementing bans of ChatGPT and other GenAI platforms. However, nearly all such bans have since been lifted; GenAI technologies have become more ubiquitous and educators have begun to understand how these tools work and to see the value of integrating AI into educational approaches. More colleges and universities are now working toward creating policies around approaches to responsible use—and many have done so already—rather than all-out bans.

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.

While initial reactions spawned excessively cautionary and prohibitive policies, conversations regarding GenAI policies in higher education have evolved to recognize the growing role of GenAI in the academic, professional, personal, and civic lives of

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 required 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.

students, faculty, staff, and administrators. Colleges and universities are now shifting their focus toward developing more nuanced and productive policies that explore these technologies' benefits in relation to higher education's objectives. Educational institutions are now exploring policies and procedures that promote responsible use, ethical considerations, and academic integrity in the context of GenAI. This includes developing guidelines for students, teachers, and researchers about appropriate AI use, integrating AI literacy across the curricula, and adapting new assessment methods that account for GenAI as part of the learning environment. This shift signals a growing understanding that **GenAI is not simply a passing trend; it is a transformative technology.** Higher education must take up proactive approaches to providing policies and resources for their campus communities.

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 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. Note, too, that the **Digital Education Council's 2024 Global AI Student Survey** reports that 86% of students report that they use AI in their academic work regularly. The incongruity between institutional policy and student use will need to be reconciled.

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. In fact, given the evolution of the role of GenAI in higher education recently, administrations really should develop policies to encourage instructors to do so.

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 (note, too, that federal and state regulations will also affect the likelihood of inconsistency). 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.

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, 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 development and research opportunities to work with GenAI

technologies. Given the anxieties that the availability of GenAI technologies have triggered among many educators, it seems judicious to formally prepare instructors—no matter their rank, teaching obligations, or contractual assignment—to engage with 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 (see Chapter 6 for more on faculty development).

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 (in the context of teaching and learning) as "understanding the fundamentals of how AI works; critically evaluating the application of AI tools in teaching, scholarship, and the management of educational priorities; and maintaining vigilance in evaluating tools and techniques to protect against bias, misuse, and misapplication of these powerful models."

Interestingly, many 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 partly 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, a pedagogical imperative and an ethical obligation exist to prepare students to understand how AI functions in those 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 simultaneously cast as posing a risk to DEI efforts and as a powerful tool for DEI.

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 some GenAI platforms do not currently charge for access to their most basic services, others do, and some provide 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. Likewise, institutions must consider access regarding disability. AI can significantly enhance accessibility for people with disabilities through tools like advanced screen readers and voice recognition, but it's critical that AI development prioritize inclusivity by addressing biases and ensuring compatibility with assistive technologies. Colleges and universities will need to take such discrepancies into account not only when developing policies but also when determining which platforms the institution will provide students and the methods for approving platforms used across the curricula (see Chapter 7 for more on compliance and platform adoption).

- *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 primarily operate in English. While efforts are underway to develop multi-lingual and other-than-English platforms, GenAI is primarily an English-speaking technology. This is because Natural Language Processing (NLP) is one of the more challenging and slower-developing areas of GenAI because of the complexities of human language and because most NLP technologies primarily rely on English training data and resources, which also highlights the dominance of English in digital content. 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 many GenAI platforms (which, notably, are growing in their 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.

Privacy

When I conduct GenAI faculty development workshops or work with administrators to develop campus AI strategies, I ask how many of them have used GenAI platforms. Often, several say that they have. I then ask, "How have you addressed privacy?" Many respond confusedly: "What are you talking about?" I then show them the privacy statements of some of the more popularly used platforms, several of which plainly describe how they collect their users' personal information and intellectual property. Simply put, when institutions ask students to use GenAI platforms—whether in the classroom or elsewhere on campus—or when administrators and researchers use

GenAI platforms, they need to be aware of the privacy issues for each platform. This is why before adopting or assigning a GenAI platform—whether the free version or the licensed version—faculty, administrators, and researchers should become familiar with its privacy statement. They should also understand some fundamental aspects of privacy and GenAI use, such as the following:

- *Data security and confidentiality:* Generally speaking, academics are alert to maintaining student privacy when it comes to personal information like addresses, grades, or other identifying information, and most are familiar with **FERPA** guiding principles about privacy. However, instructors may unintentionally ask students to disclose private information—or researchers might release intellectual property—to the platform's training data without understanding a platform's privacy policies. To access a platform or to purchase a subscription, users might be required to input sensitive personal information into a platform's registration page. Users should be informed about what data the platform collects and how the data is used. Adopters and other users should ensure that a platform provides adequate safeguards to protect privacy, including security measures to protect data in case of a data breach. Likewise, teaching faculty, staff, and students how to locate and understand privacy policies when they use platforms in all contexts is an important part of AI literacy.

- *Consent:* Adopters should ensure that they understand how GenAI tools are being used, what data is being collected, and how that data will be used. Instructors, for example, should obtain explicit consent from students, especially before using tools that require them to input personal information.

- *Transparency:* One of the most critical things we need to understand about using GenAI is the importance of transparency. Not only should we teach students to identify when and how they use GenAI in their work, but we should also be transparent about when and how all members of a campus community interact with AI tools. As educators, we need to document when and how *we* use GenAI. Such conversations help us avoid confusion regarding expectations and restrictions on GenAI use.

- *Data mining:* We should all be aware that some GenAI platforms might use data we provide for purposes beyond how we want the data used, such as for training AI models or for targeted advertising. Many platforms use prompts and other input information—like an essay or a mathematical formula—as part of their training models. That is, a platform may accumulate input information for its own purposes beyond the user's expectations. All members of the academic community should be aware of the terms of service of the platforms they use and ensure they align with ethical and privacy standards, and they should take steps to protect work from being used to train AI models without consent.

- *Intellectual Property:* Given the relative newness of GenAI platform use in higher education alongside the prolific use of these technologies outside the educational sphere, there is still significant uncertainty regarding copyright and intellectual property in the case of GenAI-generated content. Currently, for example, AI-generated images cannot be copyrighted in the US (though these laws are being revised rapidly); however, some AI-assisted writing can be. Administrators, researchers, instructors, staff, and students should be aware of the legal implications and ensure that they understand the ownership and usage rights of AI-generated work.

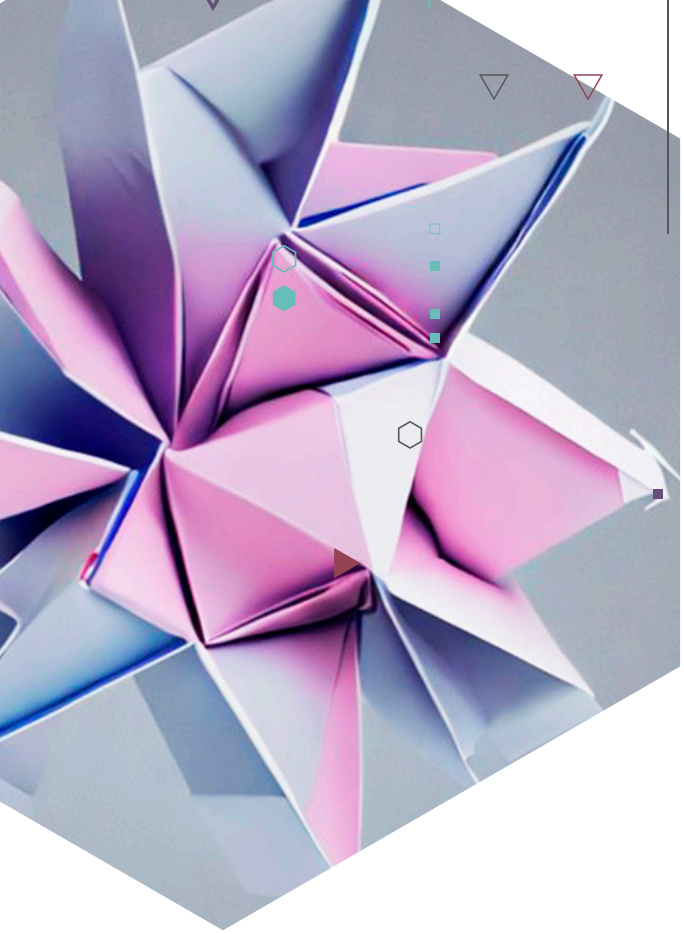

03

Artificial Intelligence for Course Instructors

Curricular Relevance (or, Why Integrate GenAI?)

As instructors begin to consider how they might integrate GenAI into their classrooms, they will need to ask themselves questions about their motives, objectives, and methods for doing so. Some of the answers to these questions may be grounded in institutional missions and practices; others will be influenced by individual instructors' educational philosophies.

Before instructors determine how they might integrate AI and GenAI into their classes, they should explore several aspects of AI integration. Such conversations should take place at all levels—across institutions and programs, and among individuals. Instructors should consider the following kinds of questions before determining how they might bring GenAI into their classrooms.

Why? All instructors should begin considering the use of GenAI in their classrooms by asking themselves, "Why do I need to bring GenAI into my classroom?" Unfortunately, the most common response to this question is one of external pressure or obligation: "I feel like my school is forcing me to." While institutional and cultural pressure may play a role in an instructor's decision to integrate GenAI, it should not be the only reason for doing so. Instructors should ask other important questions, like the following:

- Can students in my class benefit from AI literacies?
- Can hands-on experience with GenAI tools benefit and/or distract my students?
- Can GenAI promote experimentation in my curriculum?

- Can the GenAI skills that students develop in my class transfer to other areas of study, or other aspects of their lives?
- Can GenAI contribute to student efficiency?

What are the benefits? Instructors will need to carefully weigh what the benefits of bringing GenAI into their classrooms might provide their students. Consider asking these kinds of questions:

- Can GenAI tools make learning more engaging and effective for my students?
- Can GenAI help me address diverse learning styles and needs in my classroom?
- Can GenAI free up my time and resources to focus on higher-level teaching activities?
- Can GenAI help students develop essential skills for their academic, professional, civic, and/or personal lives?
- Can incorporating GenAI encourage students' deeper learning and critical thinking?
- Can GenAI unlock new possibilities for the content and methods I teach?

What are the risks? Clearly, if instructors ask about the benefits of integrating GenAI into their classrooms, they should also be attentive to the risks involved. Asking questions like these can help instructors think through potential risks:

- Can GenAI make students overly reliant on technology, hindering their ability to think independently and solve problems creatively?
- Can GenAI threaten academic integrity?
- Can GenAI use threaten my students' privacy?
- Can incorporating GenAI create an over-reliance on standardized answers and limit diverse thinking?
- Can the use of GenAI widen the equity gap between students with different levels of access to technology?
- Can GenAI integration undermine the importance of human interaction and mentorship in learning?

What is my own AI readiness? One of the most common questions I receive about integrating GenAI into the classroom is that of instructor preparedness. Because AI and GenAI have, for so long, been relegated to small segments of the academic world, many instructors have little to no experience using AI technologies in their own work and feel unprepared to teach students how to use them. It is important, then, not only for institutions to provide significant faculty development opportunities to overcome these challenges, but also for instructors themselves to account for their own AI readiness. Asking these kinds of questions can help instructors better understand their own AI readiness:

- Am I adequately prepared to use AI effectively and responsibly in my teaching?
- Would relying on AI tools risk de-skilling me as an educator?
- Do I have previous experience with AI tools?

- What is my perception of the role of AI in education?
- Do I have any personal interest in AI integration?
- Have I engaged in any training or professional development on using GenAI tools in the classroom?
- Can my institution provide me with resources and support to help me become more comfortable with GenAI use in the classroom?
- How can GenAI help me think about what comes next in how I approach my teaching?

These are some preliminary questions instructors should ask themselves when considering the use of GenAI in their classes. I encourage all colleges and universities to provide faculty development programs for all instructors through which they can engage in robust discussions about *why* instructors might consider integrating GenAI in their classes, and *how* instructors might facilitate such integrations (see Chapter 6 for more on faculty development).

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 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, as well as in social settings like social media. These kinds of inquiries can be situated in terms of integrity while simultaneously serving experiential learning objectives.

Authenticity and the Plagiarism Problem

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 accurately reflects 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, including questions about how a student develops a prompt to get a GenAI platform to generate content.

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, iThenticate, and many others 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 yet fully reliable in this capacity.

Until reliable GenAI-detection resources are available, we must 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 the 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, extending grading turnaround times. This increase in responsibilities must be addressed by individual institutions, who must 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 crediting that work to the original writer. This understanding has been so central to the core values of higher education that it has 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 more broadly, which are bound up in ideologies of capital and ownership.

- Second, the quandaries surrounding plagiarism and GenAI technologies call into question our longstanding acceptance that humans operate as autonomous 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.

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 perspectives.)

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 understanding 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, academic integrity, 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 a 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 has been 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 nowhere 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 GenAI 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 GenAI'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.

Traditional approaches to deterring student plagiarism, 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 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—some of them tailored to specific research 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—many 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. Already, GenAI platforms are beginning to provide this kind of information; but for now, their ability to do this is 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 the 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 platform 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 it 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. To this end, 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 First-Year Composition has been a near-universal course offering for all incoming college students for more than three quarters of 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 topic, 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 add 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 also necessitate 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, similar critiques have been levied throughout the history of writing about nearly every writing technology. 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 the 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 (similar to how we now understand 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

▶ **ASKING STUDENTS** to be self-reflective and to show how they moved through the *process* of writing—or the steps they took to solve a mathematical equation—redirects student learning from product to process. As instructors already know, a student's ability to explain how they reached an answer or created a final product reflects deeper learning, and possibly greater retention and transferability, than the product itself. This focus on process is not new, having gained popularity in the 1960s and 70s—but it is more pertinent than ever in the age of GenAI. Unlike human writers, GenAI can't show how it got from a prompt to a product: it doesn't prewrite, produce multiple drafts, or revise its work. Placing less value on the final product than on the student's process in assessment rubrics, then, can emphasize students' learning while also discouraging the inappropriate use of GenAI. Furthermore, actively encouraging students to find ways to use GenAI in the *process* of writing or of working through a problem—and to do so transparently—can promote the acquisition of valuable AI literacy skills.

writing-intensive courses have emerged as the central 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) and librarians, as institutions begin developing GenAI policies.

Generative AI and Reading

Much of the discussion about GenAI in higher education has thus far focused on production, and on using GenAI platforms to produce outputs. However, GenAI technologies are also altering how we read. Certainly, all of us—teachers, administrators, and students—have learned strategies for expediting our reading. We use selective reading strategies to locate specific information: tables of contents, indexes, headings, keyword searches, etc. We use summary tools like abstracts and executive summaries to determine if a text is something we need to or should read attentively for a specific purpose. We listen to audiobooks at increased speeds. We watch the movie instead of reading the book.

GenAI tools are also changing how we read and, in turn, how we and our students conduct research. Of course, GenAI does not replace attentive or close reading, but it is a valuable tool for developing more efficient reading strategies, especially in the context of reading for academic and professional purposes. Many of these strategies can benefit student engagement with assigned readings, especially in contexts where students may struggle to become invested in completing them (like textbook reading assignments in required courses). Many GenAI-driven reading strategies are inseparable from strategies for content production. For example, using GenAI platforms to generate ideas for possible writing topics may also involve using the platform to generate summaries of key concepts. **As instructors and researchers, we need to be attentive not only to the ways GenAI is changing how we write, but also to the ways it is changing how we read.** Consider, for example, these ways in which GenAI contributes to reading strategies:

- *Summary:* GenAIs can efficiently analyze and summarize just about any text and highlight its key points. They can also analyze multiple texts and identify common concepts or discrepancies among them. Imagine, for example, having twenty-five scholarly articles about a specific topic and asking a GenAI to identify their commonalities and differences, or imagine asking the platform to identify which of the articles cite a specific researcher's work and what part of that work is most frequently cited. (I sure wish I'd had such tools when I took my comprehensive exams and when I wrote my dissertation!)
- *Advanced searches:* GenAI can conduct more detailed and dynamic searches than standard search engines. Many GenAI platforms account for the context of a given

query, allowing for more accurate and relevant search results, even when dealing with complex or nuanced topics.

- *Enhanced comprehension:* GenAI can generate different types of explanations to aid comprehension. Many platforms can create simplified versions of complex texts, translate research papers into different languages, or generate data visualizations to aid readers in understanding information.

- *Personalized learning and research assistance:* GenAI can adjust to various learning styles and research needs. It provides personalized recommendations for relevant articles, suggests connections between different research areas, and offers tailored explanations or definitions based on the user's current understanding.

- *Source verification:* While many of the first GenAI platforms—particularly the free ones which brought attention to GenAI over the last few years—cannot identify the sources from which they extract information, many of the newer iterations—particularly those designed for academic work—can now identify, cite, and verify accuracy of information and help assess the credibility of sources.

By automating tedious tasks, providing deeper insights, and personalizing the research experience, **GenAI can empower researchers and students to be more efficient and effective in their pursuit of knowledge.**

Generative AI and Instructor Labor

As this resource demonstrates, the most prevalent discussions surrounding GenAI and higher education focus on teaching students to use these technologies responsibly and usefully. Certainly, these are important conversations driven by a sense of immediacy. However, it is also important for administrators and instructors to recognize that it is critical that faculty have professional development opportunities that focus on how instructors themselves can also benefit from the use of GenAI tools. The value of GenAI is in its ability to improve one's efficiency, productivity, innovation, and creativity; these kinds of improvements can greatly benefit instructors in various aspects of their working lives. However, the desire for efficiency ought not be the only reason instructors opt to use GenAI. To be of real value to us, GenAI tools need to improve the quality of the work we do in conjunction with improving efficiency. Thus, faculty development programs should help instructors understand how they can use GenAI to streamline routine tasks while also enhancing not only the quality of their teaching but also the personal satisfaction they find in their careers. While GenAI platforms are still adapting to higher education contexts, there are already many ways that instructors can integrate their use into teacherly labor:

- *Syllabus generation and revision:* GenAI can assist with drafting syllabi, providing boilerplate language for course policies, grading scales, and academic integrity statements that teachers can then revise and customize to their course-specific

WHEN I CONDUCT faculty development workshops, I often discuss one way to use GenAI that usually gets the attention of even the most resistant instructors. Before OpenAI released ChatGPT and campuses began to address student use of GenAI, I taught a class called "Digital Literacy and the Humanities" in the fall term. A few years later, I wanted to teach the course again, but I needed to update the course description to include GenAI. Rather than rewrite the description, I prompted a GenAI platform to revise the course description to include GenAI as an important part of digital literacy. The instructors I show this to are often interested in hearing how I accomplished this with the help of AI, but they are much more excited when I show them how I took the fifteen-week fall schedule and asked the platform to revise it to match a spring schedule. Instructors no longer need to make such changes manually. These are the labor-saving ways in which GenAI can be of value to instructors.

needs, including attention to specific course content, learning objectives, and assessment methods. Likewise, GenAI can provide revision suggestions for syllabus content to focus on consistency, clarity, and concision—revisions that many syllabi need. GenAI can also suggest updates to an existing syllabus for use in a new term.

- *Rubric design:* Some GenAI platforms are very good at developing rubrics for assignment development and assessment. To ensure the rubrics generated by the GenAI are useful, instructors should use specific prompts. GenAI can also help teachers revise, update, or improve previously written rubrics for use with new assignments. They can also help align the language of a rubric with that of an assignment to ensure consistency of learning objectives, assessment tasks, and rubric criteria.

- *Scaffolding:* GenAI can be a valuable tool for developing outlines, templates, and checklists to assist students as they work through complex scaffolded assignments.

- *Alternative thinking paths:* GenAI can assist instructors in developing materials to improve students' understanding of an assignment, such as guiding prompts, questions, and activities, that can stimulate critical thinking and deeper engagement.

- *Differentiated instructional materials:* GenAI can be useful in designing content such as summaries, outlines, and presentations to help learners who may benefit from alternative instructional materials. GenAI can also assist in developing different versions of assignments to accommodate diverse learning styles and needs.

- *Multimedia content:* Don't forget that GenAI can produce more than just text! There are platforms that can help produce multimedia content like images, videos, and audio.

- *Administrative tasks:* GenAI can assist with repetitive tasks like grading objective assessments (multiple choice, true or false), scheduling, email communication, and managing online learning platforms.

- *Summary:* GenAI can summarize lengthy student responses for quicker review. This can be particularly useful in early drafts or early iterations of assignment responses to help guide students in revising their responses.
- *Feedback:* Platforms can provide personalized feedback and comments on student work, highlighting strengths and areas for improvement.

04

Artificial Intelligence and Assessment

The advent of GenAI technologies has called into question the function and deployment of nearly every form of traditional academic assessment. Given the fundamental role assessment has historically played in academic processes, GenAI has created significant turmoil as institutions attempt to address pragmatic questions about assessment, including concerns about academic integrity, fairness, validity, and reliability. While colleges and universities work to address active approaches to assessment, companies developing AI assessment tools for the academic industry have also been working to develop platforms that answer many of these questions. New platforms now take up many of the pragmatic concerns instructors raise about assessment methods such as exams, group projects, portfolios, reports, standardized tests, and simulations.

Many institutions are focused on exploring strategies for maintaining the assumed integrity and value of traditional assessment methods, but we should also address the need to rethink the very function of assessment in higher education.

Why We Assess

In response to pressures for accountability, higher education has bought into the idea that assessment is a necessary and accurate reflection of educational success. Fueled by increased attention to student success, retention, and job placement as markers for institutional success, rankings, accreditation eligibility, and funding justification, assessment has become a cornerstone in higher education's scope and identity. However, the advent and broad use of GenAI initiates an opportunity not just to question

ASSESSMENT HAS ALWAYS been a feature of higher education, but it became much more prominent during the mid-to-late twentieth century. This was driven in large part by evolving theories on learning, and increased demands for accountability. Psychological frameworks like behaviorism and constructivism, alongside institutional pressures for documented student outcomes, spurred the formalization of assessment. The establishment of assessment offices, beginning with the University of Texas at Arlington in 1981, and influential reports like "Involvement in Learning" and "Time for Results," in 1984 and 1986, respectively, pushed for measurable student achievement. Pressures to account for and document student learning outcomes and program effectiveness drove institutions to take up assessment in more comprehensive and regulated ways. The 1990s and 2000s saw assessment solidify as a professional discipline, marked by professional organizations and publications, and the widespread adoption of assessment offices and data management technologies.

our traditional assessment methods but also to address a much bigger question: Does academic assessment actually work?

To rethink assessment in the age of GenAI, we should first understand the benefits of assessment we have embraced as important: to oversimplify, we justify assessment with the belief that it provides data that accurately reflects the success of both students and institutions.

For students, we employ various assessment methods to measure student progress and learning, to gauge whether students grasp and retain content and concepts, and to identify whether students meet learning objectives and fulfill learning outcomes. We use such measurements to identify students' strengths and weaknesses and track progress to adjust their learning paths (e.g., determining whether they should move on to the next level or be accepted into a major). We assess and provide feedback based on those assessments to assist student learning; feedback helps students understand what they do well and where they need improvement. Feedback also encourages self-reflexivity and critical thinking in student learning. Ideally, feedback and assessment can motivate students to engage with course materials more profoundly by identifying in which areas they need to be more attentive. Similarly, different forms of assessment require students to apply knowledge and skills in different ways, thereby reinforcing critical thinking and problem-solving skills.

At the institutional level, we use assessments to measure the quality of our programs, curricula, and faculty. We use assessment data to prove that academic programs define learning objectives and that students are meeting those objectives. We also use assessment data to verify that our programs are producing graduates who are well equipped for their careers, and we have even begun measuring student success based on career success after they have graduated. To this end, we assess and use assessment data to

identify strengths and weaknesses in our programs and curricula. We use assessment data to validate data-driven decisions about curriculum design, resource allocation, and program improvement. Because accreditation agencies often rely on data to validate and accredit colleges and universities, we assess our programs based on what data an accrediting agency declares it needs. We also use assessment approaches like course evaluations to garner data about student perceptions of classes and instructors.

Ultimately, we assess because of our faith that data can be used as evidence of success, and we place value in our assessment methods to provide that data. However, GenAI reveals many of the seams in those methods and values, and allows us to reimagine why and how we assess.

Reimagining How We Assess

Because GenAI increasingly contributes to how we teach and how students learn, higher education needs to account for its impact on the role of assessment. It should be addressed among assessment professionals and administrators at the national level to benefit from collaborative input and professional insight. It should also be addressed on campuses at the local level to contextualize assessment reform in the frameworks of institutional objectives, resources, and local needs.

As noted earlier, reimagining assessment requires seriously reconsidering some of higher education's most valued philosophies, epistemologies, and practices. It will be difficult to make changes, and if changes are made they will likely be met with resistance. Nonetheless, given shifting technological contexts, we must now have these tough conversations with an eye toward a productive future for higher education.

Reimagining assessment demands that we move away from many of the forms of assessment that we have relied on, trusted, and perhaps become complacent with. The essay, for example, has been a cornerstone not just of writing instruction but of assessing students' grasp of and ability to communicate concepts and content. In many cases, assessment methods like project-based assignments, presentations, and portfolios might better reflect student learning successes. Likewise, using GenAI technologies like dynamic assessment tools can provide adaptive feedback and customize assessment responses throughout learning pathways. Open-book and take-home exams should be rethought, given student access to GenAI tools; the assessment of such assignments should focus on how students successfully use these tools in their work rather than simply assessing the responses themselves. In many ways, this can be understood as an extension of the objectives of open-book exams, which, in part, consider how well students can use the resources available to them to complete the exam expectations. Such assessments can help evaluate how student problem-solving skills attend to using GenAI platforms in the process of responding to exam prompts.

The advent of GenAI has already fueled significant discussion in higher education about different possible approaches to assessment. While numerous other approaches are certainly percolating, authentic and personalized assessment are two ideas that have been picking up steam:

- *Authentic assessment*: Focusing on evaluating students' abilities in ways that mirror real-world contexts and challenges, authentic assessment attends to how well students can apply their knowledge and skills to perform relevant tasks beyond the classroom. Students might create products, conduct experiments, write and design proposals, or identify and solve problems in a real-world scenario. Authentic assessment emphasizes higher-order thinking, asking students to focus on problem-solving, critical thinking, innovation, and creativity. Similarly, authentic assessment encourages students to better understand from where a given discipline's knowledge evolves, and to form connections between the curriculum and its real-world applications and contexts. This can help students become more motivated and engaged with what they learn. GenAI can augment the development of dynamic simulations, collaborative projects, personalized scenarios, and interactive challenges that promote authentic assessment. These technologies can also assist in authentic assessment by analyzing student work and providing contextual, personalized feedback to encourage student learning, skills development, and deeper understanding.

- *Personalized assessment*: Moving away from universal, one-size-fits-all approaches, personalized assessment aims to tailor evaluations to meet the needs of individual students, accounting for individual learning styles and promoting individual learning paths. Personalized assessment is grounded in the idea that assessment strategies should be aligned with individual students' learning objectives, abilities, skill sets, and progress. It adjusts assessments and recommendations for learning pathways based on student performance to ensure students are appropriately challenged but not overwhelmed by evaluations that exceed their ability. The primary objective of personalized assessment is to provide students with feedback and guidance to support their future learning by calling attention to their strengths and weaknesses. Generative AI can be useful for generating adaptive assessment questions based on real-time student performance, providing personalized feedback toward the end of encouraging improvement, and creating personalized study guides, practice tests, and other learning tools. Note, however, that personalized assessment may be valuable in some contexts, but not others; for example, it may not be useful in credentialing assessments that require benchmark discipline-specific knowledge, like professional or graduate school entrance exams that require mastery of content and demonstration of particular skill sets.

Of course, rethinking assessment approaches requires addressing ethical considerations of GenAI use and its relation to assessment. **Attention to fairness and equity is more important than ever in approaches to assessment, since access to GenAI**

platforms will vary among students. Likewise, disparities in access to resources at the institutional level will affect what tools colleges and universities can provide their students, and, furthermore, resource disparities between different institutions will contribute to variations in student success in national assessments like accreditation tests or professional school admissions tests.

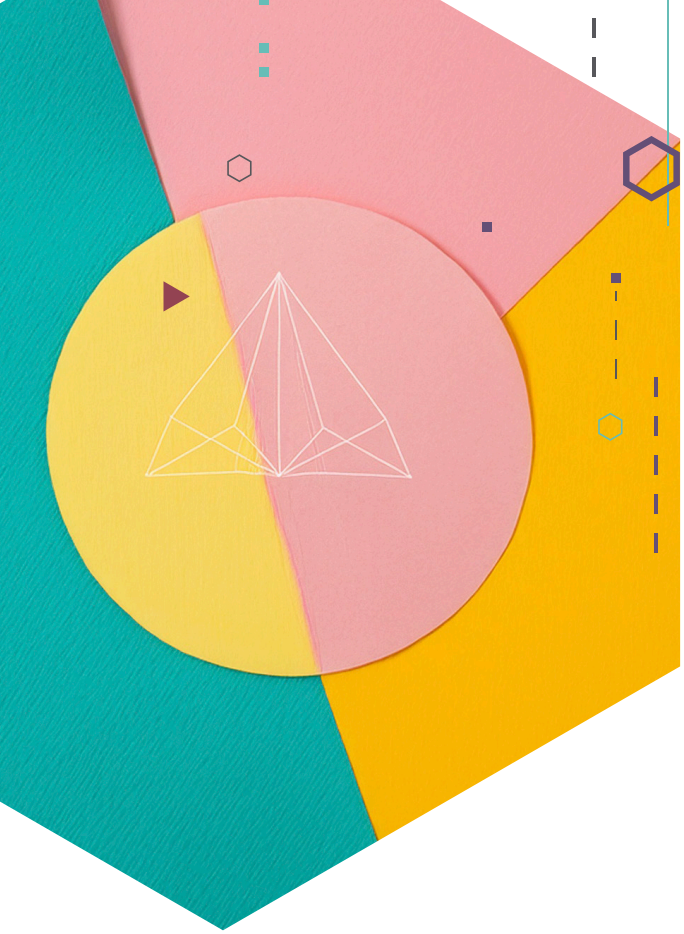

05

Artificial Intelligence and Student Career Readiness

Career readiness—the process of preparing students with the skills needed to locate, acquire, retain, and grow in their careers—is central to the missions (and expectations) of contemporary higher education institutions. This is of course particularly true for vocational, technical, and trade schools, but all postsecondary programs should teach its students the skills required to begin their careers. As more and more industries embrace the use of GenAI in their daily operations—from CEOs using GenAI to generate reports or answer emails, to graphic designers using GenAI to create product images for catalogs—higher education must now consider how the growing ubiquity of GenAI will affect how institutions prepare students for the workplace. GenAI skill is now a key component of career readiness; **understanding how GenAI interacts with specific disciplines, and with the careers that students in those disciplines 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, workplace skills, and technical skills. Academic skills include the ability to communicate clearly and effectively (particularly in written media), to perform basic research, and to perform basic computations. Workplace skills are those seen as necessary for success in the workplace; they include critical thinking, problem-solving, collaboration and teamwork, creativity, accountability, and integrity. Technical skills are job specific 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 types of skills. Any consideration about the

integration of GenAI into higher education should include significant consideration of your campus's career readiness objectives and resources. When career readiness and career placement are central to your institution's mission, representatives from career readiness programs or centers should be included in any discussion of institutional policies and curriculum development regarding GenAI.

Workplace AI Literacy

As noted above, AI literacy and readiness have become important to overall workplace readiness, and employers now identify AI readiness as one of the key skill sets they look for when hiring new college graduates. While those skills vary depending on discipline and industry, there are some generic AI literacies that employers in many professions highlight as something they would look for in a new hire. While this list is not comprehensive, these are the AI literacy skills that are often noted as important by employers:

- *Understanding of core AI concepts*: In a potential employee, organizations expect a solid understanding of basic AI terminology, such as *machine learning*, *deep learning*, *natural language processing*, *computer vision*, *robotics*, and *expert systems*.

- *Data-security awareness*: Given the value of any organization's data and the need to protect that data, employers note that awareness about data security is a top priority; understanding best practices in data security is a paramount AI readiness skill. Data-security skills include things like encryption, access controls, and data anonymization techniques, to protect sensitive information used in AI applications.

- *Data literacy*: Organizations look for potential employees who understand the different types of data the organization works with, as well as skills in data cleaning and data analysis.

- *Ethics and responsible use*: Employers understand that if their employees are going to work with the organization's data, they must do so ethically and responsibly. Employers must protect how their data is used and analyzed. More comprehensively aware organizations may also be attuned to the broader societal impacts of AI, and how the organization's data is used; employees who can attend to those broader concepts are particularly valuable.

- *Critical thinking and problem-solving*: Critical thinking and problem-solving have long been important skill sets employers look for in new hires—that has not changed. However, employers now need employees who can apply those skills in the context of AI applications. They look for potential employees who can analyze problems, identify patterns, and develop solutions using AI, including recognizing problems that emerge *because* of the use of AI.

- *Programming fundamentals*: Depending on the discipline or industry, more and more employers see value in employees who are familiar with at least one programming

language or with using platforms to develop assets (like building web pages, creating content, or designing processes). Employers especially appreciate those who can code in languages commonly used in AI development, such as Python or R. This allows new employees to contribute to AI projects and understand the technical aspects of AI systems.

- *Platforms*: Perhaps because of the velocity at which new GenAI platforms are released, employers appreciate new employees who are familiar with these platforms' various capabilities. Familiarity with multiple platforms evinces a potential employee's understanding of the practical skills needed to use each platform, demonstrates the employee's adaptability to new platforms, and allows new employees to recommend platforms and capabilities to their employers.

- *Communication and collaboration*: Because deploying GenAI in an organization's workflow may be unfamiliar to many within that organization, particularly as that organization first initiates integration, employers benefit from new hires who not only understand GenAI technologies, but who also have the ability to effectively communicate about these technologies and platforms to audiences with varying backgrounds in GenAI.

- *Domain-specific AI knowledge*: Employers look for new hires who have AI literacies along with the ability to situate their AI knowledge within their specific industry or discipline. Strong candidates should know how to transfer their general AI knowledge to specific contexts and to think critically about how AI integrations might address real-world problems.

- *Continuous learning*: GenAI technologies are evolving quickly, and employers need employees who can adapt to new innovations. Employees who can demonstrate a willingness and ability to stay abreast of the latest innovations and platforms are extremely valuable to their employers.

- *Critical evaluation*: Employers are also enthusiastic about hiring new employees who have the capacity to critically evaluate outputs from AI/GenAI platforms that an organization deploys. Employees who know how to evaluate AI/GenAI-generated information, distinguish credible from non-credible information about AI/GenAI, and identify factual claims from incorrect information about AI/GenAI within an organizational context are valuable assets to organizations.

Learning to Learn

Because AI technologies are developing at an accelerated rate, it is crucial that students be able to adapt to these frequent changes. Adaptability and skill transferability are already fundamental objectives of higher education: It's not that we want our students to learn how to write a persuasive essay about *Moby-Dick* just for the sake of learning to write a persuasive essay about *Moby-Dick*. We want students to be able to take what they

IN A DIGITAL WRITING (or content creation) class I teach, I ask students to identify an example of something in popular media that they think is "cool": a video effect, an image design, a writing strategy, a sound effect, or something else along these lines. I then ask my students first to learn how to replicate that effect, and then to teach the class how to make and use the effect. For example, let's say a student wants to learn how to create a "light saber" effect in a video. The student first figures out what resources are available to help them learn how to create the light saber effect: they locate documentation, tutorials, chat groups, or any resources they can find to help them successfully replicate it. They then have to teach the rest of the class how to replicate it, identify the sources where they found the information they needed to be able to replicate it, and describe what strategies they used to locate those sources. By the time this student has completed the assignment, not only have they learned how to recreate something cool, they have learned *how to learn* to do so.

learned from the experience of writing a persuasive essay about *Moby-Dick* and apply those skills in other contexts. The same should be true for students' use of AI.

Students need to learn how to transfer what they learned, say, from using one particular platform to how they might use another; but more importantly, **students need to learn *how to learn* about new platforms and new technologies.** That is, they need to learn how to learn.

Curricula that attend to GenAI should now include strategies for teaching students how to identify and learn about the things they don't know about GenAI. We can think of these approaches as being in line with teaching research skills, but given the speed with which new developments occur in the AI sphere, students must learn specific strategies for keeping up with these technologies. Importantly, they must also learn where and how to locate information about particular problems they may need to address. We might teach students what resources are available for conducting certain types of research in academic settings (like where to find reliable information about *Moby-Dick*); we should also teach students how to use their critical thinking and problem-solving skills, basic research know-how, and creativity to develop their own adaptable and versatile strategies for finding the information they need in order to learn the specific features of an unfamiliar platform.

Industry to Curricula

Historically, many academic disciplines have based parts—if not all—of their curricula on the credentials needed to work in a specific industry, or on preparing students for graduate or professional education to earn those credentials. In many disciplines, then, there is a direct connection between industry and curriculum; industry changes influence curriculum changes because curricula must remain relevant. Colleges and

universities need to be alert to the evolving roles of GenAI across various industries to ensure their curricula attend to those changes.

As AI skill sets evolve—and as employers are increasingly looking for these skill sets in recent graduates—institutions must develop curricula that emphasize skills for working with AI tools, analyzing AI-generated outputs, and contributing to AI development and implementation. Programs may need to develop new courses that teach these skills, and they may need to revise in-place curricula to account for these changes in industry.

One promising outcome of industry's influence on curriculum development when it comes to GenAI is a reimagined need for interdisciplinary education. Because many GenAI technologies transcend discipline-specific use or can be applied across datasets of multiple disciplines, students will benefit from learning to understand how these technologies function in multiple discipline or industry contexts; forming connections across different disciplines' problem-solving and critical thinking methodologies can enhance a student's ability to seek and provide dynamic solutions to a problem. Higher education in the era of GenAI will need to further disrupt discipline siloing and territorialism, encouraging more holistic approaches to learning. Of course, this requires rethinking many of the structural epistemologies that have divided educational institutions into distinct disciplines and programs, and imagining new ways to organize such structures will be cumbersome and counter-logistical to many, given how entrenched nearly all of our administrative structures are. Yet, GenAI integration—along with some other factors—provides opportunities to rethink our allegiances to those legacy systems. At a minimum, colleges and universities should begin to rethink program requirements; **students need more opportunities to take courses in other disciplines and to learn about other approaches to GenAI.**

Perhaps the most promising—and most problematic—result of the fast-paced evolution of industry's GenAI use is how it runs headlong into the notoriously slow process of curriculum development and approval. Colleges and universities need to reevaluate their systems for developing and altering curricula and programs to account for the high velocity at which industry changes. Approaches to efficient continuous curriculum revision in conjunction with industry trends and emerging best practices should be integrated into every institution's curriculum development and approval processes.

Such systemic upgrades should unfold hand in hand with the development of a culture of continuous faculty development. Students greatly benefit from working with instructors who have experience in or detailed knowledge of specific industries. This kind of expertise requires maintenance; instructors' understanding of industry-specific skills and standards must be kept up to date. One of the best ways to ensure faculty are aware of industry advances—including in the context of industry GenAI use—is to encourage instructors to keep up with the professional publications in their field.

In addition to changing current industries, GenAI technologies are also spawning the development of new kinds of jobs and even new industries. As these emerge, colleges and universities will need to be attuned to the ways they can educate students about the new fields so they might consider them as possible career paths. As new industries and career paths develop, higher education programs may also want to consider partnering with industry organizations in developing internship programs to provide students with practical experience. Partnering with industry organizations or creating industry advisory boards may be a practical way for institutions or programs to develop close relationships with industry professionals; this could open doors for possible guest speakers, mentorship opportunities, and industry insight in curriculum development.

By remaining in tune with industry needs and credentialing standards, colleges and universities must ensure their curricula stay relevant in order to appropriately equip their graduates for successful and fulfilling careers in their chosen fields.

IN THE FACE OF rapidly evolving industries, higher education must critically engage with the "future of work," a concept central to career readiness and student success. The transformative power of AI technologies is reshaping workplace roles, and this demands a fundamental shift in curriculum design. Coupled with the rise of remote work and the gig economy, traditional career models are dissolving, giving way to multifaceted trajectories where graduates may hold numerous jobs. This necessitates a move beyond discipline-centric education, and toward the fostering of transferable skills like critical thinking, problem-solving, and creativity, which effectively complement AI's capabilities. To prepare students for this dynamic "future of work" landscape, institutions must prioritize professional maneuverability and flexibility, potentially rethinking program structures to facilitate interdisciplinary learning. Students must also learn to effectively collaborate with AI technologies, as human–machine collaborations become increasingly prevalent. Developing lifelong learning strategies is equally vital, enabling students to adapt to constant career shifts and maintain relevance in a technologically driven world. This requires a fundamental reimagining of pedagogical approaches, ensuring graduates are not only knowledgeable but also agile and adaptable in the face of unprecedented change.

06 Faculty Development

It has become evident that a key component of every institution's approach to the role of GenAI technologies on its campus must include a robust faculty development program. Because of the relative newness of this need, leadership on many campuses is grappling with how to provide these programs. They must identify what resources are needed to support such efforts, and where those resources will come from. They must also consider how to provide expertise—especially when there are no experts already on campus—and how to coordinate faculty development programs with the development of campus-wide GenAI policy and infrastructure.

Given the speed at which GenAI technologies are becoming prevalent in higher education, faculty development is now a top-tier need on every campus. Unfortunately, because of the complexities in establishing GenAI faculty development programs and the lack of resources, at many colleges and universities such responsibilities have fallen to individual faculty by default. In the current climate, it is irresponsible for institutions to submit to an unstructured "figure it out for yourself" approach, and **faculty development must become a core mission for every institution in higher education.**

Faculty AI Readiness

It has become common practice in industry for consultants and strategists to develop GenAI approaches by assessing an organization's AI readiness. The consultant needs to understand an organization's specific needs and capabilities in order to tailor their recommendations appropriately. Often, this is accomplished by way of an AI Readiness

Survey (AIRS), a method for assessing how prepared an organization is to adopt and implement AI technologies. In my own work as an AI strategist, I use an AIRS as the first step in developing an AI strategy for each company I work with, and I have begun to transfer this assessment strategy into different academic contexts.

In the context of higher education, AIRS data can provide insights into several aspects of an institution's AI readiness:

- Faculty AI awareness, perception, attitudes, and experience
- Faculty needs for the development of AI/GenAI curricula
- Student AI literacy and skills
- Concerns about data privacy and security practices
- Faculty and student positions about responsible use and academic, professional, civic, and personal integrity
- Existing AI infrastructure and resource availability
- Need or desire for faculty development and training programs
- Discipline-specific AI research capabilities and initiatives
- Leadership's attitudes toward AI adoption
- Institutional culture around and perception of AI
- Existing (or lacking) AI applications in teaching, learning, and administration

Of course, as with any survey data, the results will be influenced by the design of the survey. Therefore, when establishing an AIRS with the goal of informing faculty development programs or increasing administrative awareness of the campus climate, it is important to involve campus personnel with expertise in the design and assessment of surveys; if this expertise is unavailable on campus, an external consultant should be brought in to assist.

Faculty Needs

Despite widespread early resistance, instructors are increasingly recognizing the potential of GenAI to transform their teaching practices, but they also acknowledge the need for support and resources to integrate these technologies effectively. One of their primary needs is comprehensive training that goes beyond basic awareness and delves into the practical applications of AI in pedagogy. In most of the faculty development workshops I facilitate, instructors regularly ask, "How do I do that?" Instructors across disciplines express a need for hands-on experience with GenAI tools, guidance in designing GenAI-integrated activities and assignments, and strategies for teaching their students to use these technologies responsibly.

Likewise, instructors at nearly all of the institutions I work with—whether two-year vocational colleges, Research-1 universities, liberal arts colleges, or higher learning institutions of virtually any other kind—regularly express the need for access to robust technological infrastructure and support. This includes reliable Internet connectivity,

access to appropriate AI software and platforms, and technical support to troubleshoot issues and ensure the smooth integration of AI tools in the classroom. Concurrently, they want direction in platform selection, an issue that requires guidance about cost, compliance (more about this in Chapter 7), and accessibility, as well as platform capability and applicability. These needs extend beyond an institution simply providing a campus-wide recommendation or license for a generic platform.

Faculty also want clear guidelines and policies from their institutions regarding the responsible use of AI in teaching and assessment; these are complex, often unclear, and equivocal areas around which to develop policies and practices. This is also murky territory for administrators to tread since, on the one hand, instructors want institutional policies to guide them, thus deferring difficult decisions to administration, but, on the other, institution-wide policies risk restricting discipline-specific practices and instructor autonomy.

Instructors also express a need for collaborative environments on campus where they can share experiences, challenges, and best practices related to GenAI integration. That is, they want opportunities to talk with and learn from their colleagues. Instructors correctly see value in learning communities, workshops, and mentoring opportunities that allow them to connect with colleagues and learn from each other's experiences. Institutional leadership should provide resources for facilitating these kinds of collaborative opportunities. Such initiatives can be instituted and maintained by centers for teaching excellence on campuses that have these entities. Many institutions have already developed such programs, and cross-institutional collaborations can provide useful examples to other institutions hoping to do the same. Likewise, learning communities can also be developed in coordination with other faculty professional development programs offered.

Professional Development Programs

Faculty development programs need to both address immediate needs and anticipate long-term needs. Colleges and universities should provide professional development programs to address immediate global concerns across the entire campus. While no one-size-fits-all approach to faculty development can be exhaustive, institution-wide programs can invigorate productive discussions about discipline- or context-specific considerations and identify areas for further development.

Faculty development programs should empower instructors with the skills and confidence to integrate GenAI into their classrooms in ways that support their teaching objectives. This cannot be accomplished with individual or infrequently offered workshops or discussions. Rather, institutions need long-term programs offered in recurring cycles to encourage faculty participation and account for faculty turnover.

Such programs should begin with the introduction of foundational concepts and ethical considerations, and culminate in the development and implementation of individual, context-specific AI-enhanced teaching practices. Early in the program, faculty should be introduced to the fundamentals of AI and GenAI, and explore their capabilities, limitations, and ethical implications. These kinds of foundational introductions lay the groundwork for teacherly AI literacies and contribute to an instructor's confidence in their ability to understand and integrate GenAI technologies that were previously unfamiliar and, perhaps, daunting. Like any effective professional development mechanism, such a program should provide hands-on workshops to encourage instructors to experiment with various AI tools and platforms, gaining practical experience in applications relevant to their particular disciplines.

Building on this foundation, development programs should then take up pedagogical integration. Instructors should have the opportunity to collaborate with one another in the design of GenAI-integrated learning activities. Other classroom applications of GenAI should be covered, such as the design of syllabi, rubrics, and assignments, as well as the potential for personalized learning and adaptive feedback. Instructors should be provided with opportunities to discuss the complexities of topics like assessment and authenticity, and how attitudes and concepts in higher learning are adjusting in light of GenAI. Strategies should be explored for fostering AI literacy among students and for empowering them to become responsible and ethical users of these technologies. Akin to student learning pathways, faculty development pathways encourage engagement and confidence, and they instigate collaboration and community among participants. They also signal institutional support, boosting faculty investment. Long-term professional development is about more than just keeping up with the latest technologies; it is about harnessing the opportunities that GenAI provides and improving the quality of the educational experience.

AI Literacy for Instructors

Because GenAI technologies are as new to instructors as they are to students, instructors should learn AI literacies to better understand how these platforms function both generally and in the contexts of their professional lives. Professional development programs should introduce instructors to fundamental AI literacy skills to boost their understanding of GenAI and improve their confidence in integrating these tools into their work. Given instructors' rampant sense of imposter syndrome regarding their own AI literacies and their abilities to teach with or about these technologies, professional development programs should begin by addressing subjects such as these:

- *Basic AI concepts and definitions*: Instructors should be able to define terms like AI and GenAI. Other key concepts include machine learning, deep learning, large language models, and natural language processing; it is also important to understand how these technologies differ and/or relate to one another.

- *How GenAI works*: Instructors should understand how AI systems learn from data, make predictions, and perform tasks, as well as how platforms make mistakes, like delivering hallucinations.
- *Various types of AI*: Instructors should be familiar with different types of AI, such as natural language processing, computer vision, expert systems, and robotics.
- *Applications of AI in education*: Instructors should be aware of at least some of the ways that GenAI tools are being used in education, such as personalized learning platforms, automated grading, and chatbots, to name but a few.
- *Potential benefits and challenges of AI in education*: Instructors should be exposed to a balanced perspective on the potential benefits, challenges, and risks of AI integration, including ethical considerations. They should also have opportunities to share their thoughts about these matters.
- *Data privacy and security*: Instructors should understand the importance of protecting student privacy and ensuring responsible data use in AI applications.
- *Bias and fairness in AI*: Instructors should be aware of how bias can be embedded in AI algorithms and the importance of promoting fairness and equity.
- *Critical evaluation of AI*: As teachers expand their AI literacies and become more comfortable working with these tools, they should also be encouraged to critically evaluate AI tools and resources, considering their limitations as well as their possibilities.
- *AI literacy for students*: As teachers expand their AI literacies, they should be encouraged to develop strategies for teaching students about AI concepts, ethical considerations, and responsible use.

Labor and Job Security

As you might imagine, in addition to the excitement about integrating GenAI technologies into the classroom, there is also significant anxiety around doing so. Much of the professional angst regarding GenAI in higher education builds upon the serious questions these technologies have illuminated about our fundamental teacherly beliefs and practices. Equally as concerning, introducing GenAI technologies has exacerbated concerns among teachers regarding labor and job security. As popular media has rightly asked how AI technologies might eliminate or change jobs in many industries, instructors can't help but wonder whether these technologies might also replace or drastically alter higher education labor. Labor experts acknowledge that AI will inevitably affect employment across nearly every industry. **A recent report by the World Economic Forum** says that approximately 41% of employers report plans to downsize their workforce as AI automates specific tasks. Further, the report explains that 77% of

I RECENTLY DELIVERED a keynote address about emerging technologies and higher education at a large public university. During the Q&A, a faculty member posed this question: "I am sixty-five years old. I have been teaching here for thirty-five years. I love what I do. I have no intention of doing anything you talk about, and I think that what you are talking about will be detrimental to teaching. What do you want me to do?" My response was inappropriate, poorly thought out, and wrong: "With all due respect, and while I acknowledge that your students have benefited from your expertise throughout your career, what I want is for you to retire." My response has haunted me since I spoke it. I further explained that technological changes require cultural shifts and time to become familiar and accepted, and that the only way to expedite change is to bring in faculty who see the value in new technologies. I realized, however, not only that my response was disrespectful but that it established a culture of obsolescence. I now actively work away from such positions and speak with faculty and administrators about avoiding cultures of obsolescence. Instructors who choose not to engage with GenAI technologies deserve respect. Their experience, creative curricula, and mentorship are invaluable. We must foster an environment that honors diverse teaching styles and expertise, recognizing that good educators remain effective regardless of technological adoption. Marginalizing resistant voices devalues their important contributions and stifles diverse pedagogical approaches.

the one hundred companies surveyed plan to reskill and upskill their existing workers between 2025 and 2030 to better integrate AI collaboration. Information like this—particularly when reported alongside popular media's fearmongering about AI—inevitably leads teachers to wonder how the use of GenAI technologies might affect higher education labor. While I acknowledge that I have promoted the value of these technologies' ability to augment and enhance academic labor throughout this resource and in many other venues, it is also important to take the concerns about teacher labor seriously as colleges and universities develop policies and processes for GenAI integration.

Like workers in other professional labor forces, instructors are concerned about job displacement, with AI potentially automating routine tasks. Teachers are also concerned about the changing nature of teaching, and as I have repeated throughout this resource, AI and GenAI technologies certainly expose a need to rethink many of our fundamental understandings of the purpose of higher education and our methods for fulfilling those purposes. Many instructors find doing so an affront to their professional beliefs, and colleges and universities must find ways to negotiate that delicate terrain.

Simultaneously, the integration of GenAI technologies into higher education has also triggered instructor anxiety about their AI literacies and about their need for professional development. While professional development is addressed in previous sections of this chapter, it is important to be aware that those same concerns also contribute to apprehensions regarding labor and increased workload in order for instructors to participate in professional development programs. Instructors regularly explain that

they are already busy—if not overworked—and that adding professional development (whether deemed necessary or not) to the litany of demands on their time exacerbates their unease about bringing these technologies into their classes. Not receiving professional development support, on the other hand, potentially leads to anxiety about keeping up with technological advancements, individuals' perception of their own professional relevance, and the unfolding redefinitions of expectations from their teaching.

As institutions integrate AI and GenAI technologies across their curricula and navigate the rapidly changing landscape of AI in education, it's crucial to recognize that AI is likely to transform the labor and function of instructors. Administrators need to account for how these changes will benefit institutions, certainly, but must also be alert to the impact they are likely to have on the lives of their instructors.

07 Infrastructure

It is one thing to acknowledge that colleges and universities need to find ways to integrate AI technologies across their institutions; it is another for them to provide the resources and infrastructure to do so effectively. Because of the velocity at which AI technologies rapidly moved from cursory to necessary across higher education, many institutions are scrambling to figure out not only how to integrate these technologies, but also how to provide the resources needed to do so in more than superficial ways. Ranging from costs for hardware and platform licenses to upskilling teachers, researchers, administrators, and staff, incorporating these technologies requires the commitment of resources, which many institutions may not be able to do easily.

Certainly, integrating AI technologies at any college or university will benefit that campus's community and provide opportunities to transform research and teaching, but doing so requires robust support and resources and equitable distribution.

Technological Support

At the most basic level, institutions will need to ensure that their community has equitable access to high-speed Internet that can handle the large bandwidths that many AI platforms require. They will need to provide consistent Internet access across campus, guaranteeing access in classrooms, offices, libraries, labs, communal spaces, and housing areas. In tandem with Internet access, institutions should provide access to computers with large computing and processing capacities. This may require re-investment in computer labs (which many campuses have moved away from in recent years due to cost and the proliferation of students having their own devices), but it may also require

rethinking access policies. Many campuses and even specific programs require students to have their own computing devices, like laptops, that meet specific processing capabilities. Those who oversee technology policies on campus or in specific programs may need to reevaluate their current requirements and revise access policies based on new AI/GenAI needs. In doing so, they will need to consider the financial burden that any new policy may place on students and how the new requirements might interact with financial aid, fellowships, and scholarships. They will also need to identify from where resources might be harnessed for developing or redesigning on-campus resources, such as computer labs to provide more robust remote access capabilities. And, of course, institutions will have to consider the cost of providing licenses and subscriptions for AI/GenAI platforms, development platforms, and data access.

This also requires institutions to determine how platforms will be selected. Many colleges and universities have begun to purchase licenses for all students on campus to have access to platforms that are likely to be of value to large swaths of a campus community. Many colleges and universities now provide students with access to general writing or content creation platforms like Gemini, ChatGPT, Copilot, and Grammarly. Others provide students across campus with licenses to platforms for working with images, videos, audio, and design, turning to campus-wide licensing of tools like Adobe's Creative Cloud through Adobe's Creative Campus program. However, there are numerous contexts across campus that will benefit from more discipline-specific platforms. Administrators will need to address how to account for smaller-sized platform needs to figure out how to provide those platforms in cost-effective ways. For example, we may need to begin to assign platform registration and licensing in the same way we assign textbooks, deferring costs to students. Such approaches will need to account for coordination with financial aid, scholarships, and fellowships, as well as coordination with campus bookstores for providing discounts and methods for adoption.

Institutions will need to develop durable data management and security approaches and policies. Matters of data security will need to be upgraded at many institutions, particularly as instructors, researchers, and administrators begin working with large datasets in ways they previously had not. This may also require institutions to rethink their data storage approaches, potentially needing to upgrade local servers or cloud storage access. Institutions will need to develop more rigorous policies regarding the safekeeping of data and the assurance of compliance with relevant regulations (more on this below).

Likewise, institutions will need to provide comprehensive technical support. This may require upskilling of technical support personnel or hiring new support personnel with expertise in AI technologies and platforms. Support staff should also be able to provide basic training through workshops, webinars, tutorials, online resources, and FAQs to better prepare instructors, researchers, and staff for using these technologies. Such approaches should be integrated into larger training and development support

programs; lifelong learning approaches for individual professional development should be encouraged, but individual approaches cannot be the only form of support.

Structured professional development programs are necessary to establish a campus-wide learning community, consistently address campus policies and approaches, and ensure that faculty development parallels the available resources and campus capabilities.

Training and Development

In addition to the provision of hardware, software, and technical support, faculty development opportunities are crucial to the successful onboarding of AI technologies in higher education. Administrators know all too well that providing such programs requires resources, and new strategies may need to be developed for reallocating or acquiring the resources necessary for providing comprehensive professional development programs to support faculty, researchers, and staff in integrating AI/GenAI technologies to advance the institution's objectives.

Administrators will need to consider the return-on-investment potential of the various possible approaches. They should consider, for example, whether to prioritize upskilling current members of the campus community or to hire new faculty or staff with AI expertise. In particular, in the long run, it might be more cost effective to hire new personnel in technological support positions who can facilitate AI upskilling and development programs than it would be to have current staff develop the expertise needed to do so. In all likelihood—and this trend is already unfolding nationwide—campus IT divisions will need to invest in hiring AI specialists to oversee AI integration and to develop support networks on campus.

The kinds of professional development support that campuses should offer include things like workshops, training sessions, and online resources to help faculty understand AI concepts, learn how to use AI tools in their teaching, and develop AI-related curricula (see Chapter 6 for more on faculty development). However, what may be just as important during the early phases of integration is the presentation of professional development support for everyone in the campus community as an integral part of campus culture. AI integration will be more widely accepted when members of the campus community are actively involved. Thus, administrators should provide all campus community members opportunities to contribute to discussions about AI use on campus, including decision- and policy-making conversations, and ensure that those conversations are taken seriously by those who control campus resources.

While much of this guide focuses on integrating GenAI in teaching contexts, administrators must also consider ways to support staff development. AI technologies can be beneficial to those who do the work of running colleges and universities; beyond technological support staff, nearly all members of the higher education labor force should

be provided with training and platform access. It makes little sense to provide professional development to instructors without providing resources to those who support those instructors' teaching. AI technologies can benefit nearly all campus community members, from academic advisors and career counselors to residence life staff, from library personnel to media services, and from human resources to security services.

Compliance

As AI integration continues to become more prevalent across higher education, administrators will face many compliance issues. These range from monitoring individual use compliance on campus to an institution's compliance with state and federal regulations to compliance with various funding agencies' policies. The intricate landscape unfolding as a result of the growing ubiquity of AI in higher education requires the careful attention of administrators at all levels of an organization's leadership. Administrators must take up data privacy, academic integrity, bias, accessibility, equity, intellectual property, copyright, transparency, alignment with institutional missions and policies, and compliance with state and federal regulations.

Institutions need to be mindful of data privacy policies and safeguards of the platforms that are deployed on campus, whether for research, teaching, or support. To that end, colleges and universities must develop information security risk assessment processes that can account for campus-wide, course-specific, and research-related deployment of various AI platforms. Such assessments can be time-consuming, so administrators should develop efficient policies for individuals to submit requests for assessments of platforms they wish to use. Such processes should account for time frames of need; for example, a teacher designing a course for an upcoming semester will need timely platform approval in order to use the platform in question with their students.

Administrators will also need to align their data security policies with external policies, like FERPA, the General Data Protection Regulation, the Americans with Disabilities Act, copyright and trademark laws, and sector-specific regulations where applicable, among many others. Administrators will also need to develop policies for the ways an institution's data is securely stored, particularly when cloud-based storage is used. Storage access, accessibility of access controls, data-sharing agreements, data retention, data breach monitoring and notification, and vendor security practices and audits need to be developed.

Because AI in higher education implicates intellectual property and copyright, compliance policies require careful attention to information ownership, focusing specifically on protecting the intellectual property of students and researchers. Because GenAI technologies specifically raise questions about content authorship, colleges and universities will need to address concerns about who owns content generated by GenAI platforms, particularly when the content is based on data acquired by employees of the institution. Likewise, policies regarding platform adoption will need to consider

whether the platforms capture input data, like prompts or metrics, for use in their own training models. Policies need to specifically address how various platforms retain and use student data when students are using such platforms for assigned work.

One of the difficulties in both developing compliance policies and monitoring compliance is that the AI technologies for which these must be developed are not stagnant—they continue to evolve at a rapid pace. New platforms and capabilities emerge almost weekly, and, often, other regulatory policies change to account for these. Thus, **institutional policies need to be organic, fluid, and adaptable.** Colleges and universities must regularly review and update compliance policies to attend to changes in technological development, licensing agreements, campus use, sector-specific policy adjustments, and other regulatory changes.

Sustainability

Much of this resource treats the integration of AI technologies in higher education as something new, which implies that many campuses may be in the early stages of integration or may still be trying to figure out whether they should bring AI to their campus at all. However, it is more accurate to write that colleges and universities have already initiated integration processes and are now exploring ways to build, support, and expand early initiatives into longer-term efforts that support their institutions' missions. That is, it is now crucial for administrators to take up the long-term sustainability of AI integration. Following the initial excitement and anxiety about AI technologies, it's time to accept that they will be a permanent core component of higher education. Administrators must now find ways to ensure that these tools remain accessible, that professional development programs expand, and that teachers continue to develop new curricula to provide students with robust AI-integrated learning experiences. This

▶ **WHEN IT COMES TO** AI integration on any campus, public perception is crucial, given the often polarized public perception of these technologies. Institutions have a unique opportunity to provide informed perspectives, clarifying their visions, plans, and the tangible benefits of AI deployment; transparent discussions about research advancements, efficiency improvements, and curriculum innovations build public understanding and trust. Furthermore, fostering a culture that welcomes feedback from those inside and outside the campus community is critical. This engagement supports community investment, encourages dialogue, and promotes inclusivity, particularly within the hierarchical structures of higher education. Public anxieties surrounding AI must be addressed. Institutions should proactively communicate their policies and the processes behind their development, fostering confidence and establishing clear touchpoints. By actively shaping the narrative and fostering open communication, universities can effectively navigate the complexities of AI integration and build a supportive environment for its implementation.

means that administrators need to develop strategic procedures to address long-term, sustainable approaches to financial, structural, and curricular provision.

Financial sustainability may be the most difficult aspect of AI sustainability, given the limited and often strained budgets many campuses face. Acquiring new sources of financial support or reallocating resources from other campus budgets may prove difficult. Not only will colleges and universities need to continue to support up-front investments in hardware, platform licenses, data storage services, and other expenses, but budgets should also include plans for ongoing maintenance, upgrades, replacements, new license agreements, cloud computing costs, personnel, faculty development, and technical support.

As noted several times in this resource, faculty and staff development for AI integration requires significant investment and should entail much more than a one-time introduction. Appropriate professional development requires commitment to and investments in sustainable and expanding programs; this includes training new personnel, paying faculty and staff for continuing education efforts, and the creation of training resources.

08 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 resist the use of AI in education. Many wish to ban these technologies; many 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.

To be clear, I do not want to whitewash the story of the Luddites. The mechanization of the textile industry did result in the loss of some jobs; it also caused significant violence and loss of life, as those opposed to mechanization took up a campaign of violence to destroy mechanized textile machines, and the industry responded with violent strategies not just to repel the Luddites' destructive acts but to impose power and silence any action that questioned their industry. That is, I recognize the danger and inaccuracy in characterizing the Luddite narrative as having an "OK" ending in the long run.

The story of the Luddites does, however, make 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. 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

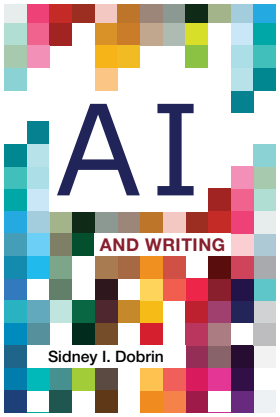

AI and Writing provides a road map for students and instructors who wish to talk about GenAI programs such as ChatGPT but who don't know where to begin. Writing prompts, discussion questions, and provocative short readings guide students in their consideration and use of GenAI.

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.

Keep an eye out for *AI and Writing*, Second Edition, forthcoming in early 2026.

Click here to request an examination copy of the first edition 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

SID DOBRIN, chair of the University of Florida's English department, has become one of the world's most sought-after academic experts on GenAI. He is the founding director of the Trace Innovation Initiative at UF and has been named a digital thought leader by Adobe. He serves as a member of the Florida Institute for National Security, part of the University of Florida's AI Initiative, serves as a member of the Florida AI Learning Consortium (FALCON) Steering Committee, and as a member of UF's AI Blue Sky Task Force. He is the author and editor of numerous books and articles, including *AI and Writing* and the forthcoming collection *AI and the Humanities*. His current research project examines enduring questions motivated by AI. In addition to his work in digital humanities, he is a prolific writer about environmental humanities, specifically focused on blue ecocriticism and blue humanities.

Sidney I. Dobrin is available to consult 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 identify 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.

broadview press
www.broadviewpress.com

ISBN 978-1-55481-704-7