

EDUCAUSE

REVIEW



ARTIFICIAL INTELLIGENCE

Where Are We Now?

SPECIAL REPORT

Thank You

EDUCAUSE thanks the sponsors
of this Special Report on artificial intelligence.



EMPOWER yourself and your organization.

The community, tools, and resources
you need to achieve your mission.

educause.edu/discovermembership

Contents

JUNE 1, 2022

ARTIFICIAL INTELLIGENCE WHERE ARE WE NOW?

The *EDUCAUSE Review* Special Report on artificial intelligence (AI) curates and contextualizes the best of AI content from EDUCAUSE. Articles will help higher education executives and senior IT leaders understand where we are now with the use of AI and its applications in higher education, its promises and perils, the ethical implications, and its role in ensuring student success.

4 Introduction: Moving Beyond Pockets of Excellence in Higher Ed AI

John O'Brien

Is AI changing the world? Is it transforming higher education as we know it? Or is it working quietly in the background?

6 Artificial Intelligence: 7 Questions (and Answers) [infographic]

8 Artificial Intelligence: Threat or Opportunity?

Brian Fleming

Half of Americans surveyed said they trusted higher education to build, manage, and govern artificial intelligence. We should lean into this finding.

12 Artificial Intelligence and Ethical Accountability

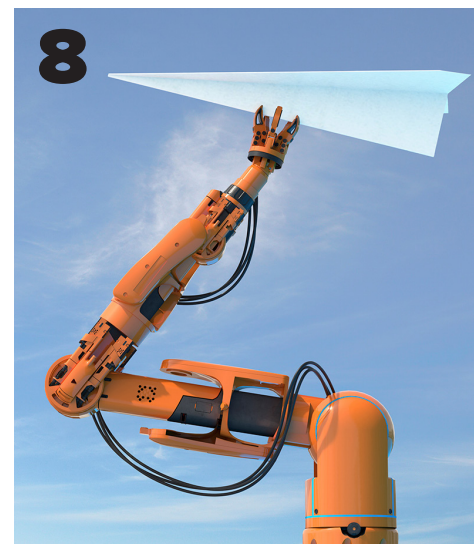
Linda Fisher Thornton

When selecting AI-enabled programs for hiring, teaching, testing, or student support, higher education institutions should weigh the cost savings and efficiency of the programs against any possible risks to fairness, privacy, bias, and public safety.

18 EDUCAUSE QuickPoll Results: Artificial Intelligence Use in Higher Education

D. Christopher Brooks

Despite the visibility and energy surrounding artificial intelligence, higher education is taking measured steps toward incorporating this technology.



26 How Higher Ed CIOs Can Leverage AI

Toby Jackson, Mainstay

28 Artificial Intelligence in Higher Education: Applications, Promise and Perils, and Ethical Questions

Elana Zeide

What are the benefits and challenges of using artificial intelligence to promote student success, improve retention, streamline enrollment, and better manage resources in higher education?

40 3 Ways AI Can Help Students with Disabilities

Artificial intelligence can improve accessibility and ensure that students with disabilities have access to rich learning opportunities.

42 Combining AI with Human Oversight: 3 Benefits for Online Proctoring

Jordan Adair, Honorlock

44 Preventing a Winter of Disillusionment: Artificial Intelligence and Human Intelligence in Student Success

Linda Baer, Amanda Hagman, and David Kil

Using artificial intelligence to better inform human intelligence, higher education can prevent a winter of disillusionment and ensure tangible student success outcomes.

54 For Higher Ed, Chatbots Need to Be Smart

Mary Frances Coryell, Ivy.ai

56 No Cheat Code: Exploring the Intersection of Empathy and Technology

Tara Hughes

When technology is implemented thoughtfully and empathetically, the impact can be profound.

Cover: VS148/Shutterstock

Why IT Matters to Higher Education
EDUCAUSE
REVIEW

EDITOR-IN-CHIEF
D. Teddy Diggs

DIGITAL MANAGING EDITOR
Karen Henry

MULTIMEDIA PRODUCER
Gerry Bayne

SENIOR EDITOR
Gregory Dobbins

ADVERTISING
Carolyn Colman, Senior Manager, Membership Development
Che DeMarco, Membership Development Specialist

DESIGN AND PRODUCTION
Karen Henry
D. Teddy Diggs

EDUCAUSE EXECUTIVE OFFICERS
John O'Brien, President and CEO
Eden Dahlstrom, Vice President, Professional Development
Susan Grajek, Vice President, Partnerships, Communities and Research
Mairéad Martin, Vice President and Chief Information Officer (CIO)
Nicole McWhirter, Chief Planning Officer
Catherine Yang, Vice President, Digital Communications and Content

EDUCAUSE
UNCOMMON THINKING
FOR THE COMMON GOOD

EDUCAUSE Review is the association's award-winning flagship publication for the higher education IT community. The magazine takes a broad look at current developments and trends in information technology, how they may affect the college/university as an institution, and what these mean for higher education and society. It also offers practical advice and guidance on using information technology in higher education.

EDUCAUSE and *EDUCAUSE Review* are registered trademarks. Copyright © 2022 by EDUCAUSE. Statements of fact or opinion are made on the responsibility of the authors alone and do not imply an opinion on the part of the EDUCAUSE Board of Directors, staff, or members. For more information about copyright, see <https://www.educause.edu/copyright>.

4845 Pearl East Circle, Suite 118
PMB 43761
Boulder, CO 80301
phone: 303-449-4430
fax: 303-440-0461
info@educause.edu
<https://www.educause.edu>

For advertising information, e-mail advertising@educause.edu or call 303-544-5664.

EDUCAUSE Review is available online at <https://er.educause.edu>.

Introduction: Moving Beyond Pockets of Excellence in Higher Ed AI

By John O'Brien

F


or a while now, I've often shared some of the more dramatic statistics and predictions about artificial intelligence, like the claim that nearly 30 percent of consumers couldn't say if their last customer service exchange was with a human or a bot and the prediction that by 2020 "the average person will have more conversations with bots than with their spouse."¹ However, as is often the case, perhaps the more important number is the more nuanced (and less dramatic) one.

As far back as 2017, reports have surfaced about the surprising proportion—sometimes one-half, sometimes two-thirds—of people using AI in their daily lives who do not believe they are doing so.² It's one thing to nod to the rapid growth of AI adoption, but it's something else entirely to track the degree to which a growing reliance on artificial intelligence is happening under our noses without our even knowing it. Is AI changing the world? Is it transforming higher education as we know it? Or is it working quietly in the background?

This *EDUCAUSE Review* special report highlights promising practices in the use of AI in higher education. AI can make a significant difference: for example, incorporated in ways to better serve students with disabilities and used (adaptive learning) to improve student

success with demonstrable results. AI can give students complementary opportunities to dig into course materials, and when it's working well, it can free faculty from transactional tasks and unleash them to interact, as only humans can, with their students.

However, as noted in the EDUCAUSE QuickPoll on AI, "Current use of AI is a mile wide and an inch deep." The examples remind me of how we used to describe campus initiatives that worked in a few departments but weren't adopted on a widespread basis. We smiled optimistically and observed that we were seeing "pockets of excellence." The gap between where we are and where we hope to be is perhaps most obvious in research from McKinsey and Company, a global management consulting firm. AI could, the company reports, reduce teachers' workloads by 20–40 percent and



ALESIAKAN/SHUTTERSTOCK



cut prep time and administrative time in half.³ The data point is inspirational. Yet I found myself thinking that AI in higher education will remain an unrealized hope until those in the field, not consulting companies, are the ones making the case for increased adoption.

When will we find traction for AI in higher education? When faculty are clamoring for more. Trust me, when the value proposition for AI is the liberation of faculty from the drudgery of transactional interactions with students ("When is the test?" "Where is the lab?" "When is the paper due?"), faculty will stand in line to sign up. Until then, higher education will continue to be in an adolescent phase in the use of AI.

Meanwhile, one of the most important ways to build the case for AI in the field is to delineate the evidence of impact shown by

studies exploring how AI improves student outcomes. Similarly, research is needed to examine the unintended pitfalls of AI—with a deliberate focus on what ethical AI looks like, with hard work at the point of design (and not as remediation done after an AI solution has launched).⁴

You'll see these themes and many others woven into this special report on AI. And the time is decidedly right. Higher education professionals need to be aware of the directions in which AI is moving within our world. There is little question that what we do (and decide not to do) at this early stage will have important implications for the decades that follow.

Bump. Bump.

I know it's time for me to stop reflecting and writing about the potential of artificial intelligence when my AI-driven portable vacuum cleaner is repeatedly bumping against my office door. It has mapped our floorplan, and it knows I am in here working. Meanwhile, it's working quietly (mostly) in the background to transform my dirty floors. ■

Notes

1. "Bot.Me: A Revolutionary Partnership," PwC Consumer Intelligence Series, 2017; Heather Pemberton Levy, "Gartner Predicts a Virtual World of Exponential Change," Gartner (website), October 18, 2016. For more forecasts from a few years ago, see Gil Press, "AI by the Numbers: 33 Facts and Forecasts about Chatbots and Voice Assistants," *Forbes*, May 15, 2017.
2. Shep Hyken, "Half of People Who Encounter Artificial Intelligence Don't Even Realize It," *Forbes*, June 10, 2017; Eileen Brown, "Two Out of Three Consumers Don't Realize They're Using AI," *ZDNet*, March 16, 2017.
3. Jake Bryant, Christine Heitz, Saurabh Sanghvi, and Dilip Wagle, "How Artificial Intelligence Will Impact K-12 Teachers," McKinsey & Company (website), January 14, 2020.
4. For more on the ethics of AI, see John O'Brien, "Digital Ethics in Higher Education: 2020," *EDUCAUSE Review*, May 18, 2020.

John O'Brien is President and CEO of EDUCAUSE.

© 2022 John O'Brien. The text of this work is licensed under a [Creative Commons BY-NC-ND 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/).

This article was originally published in *EDUCAUSE Review* on [June 1, 2022](https://www.educause.edu/pressroom/2022/06/john-obrien-introduction-moving-beyond-pockets-of-excellence-in-higher-ed-ai).

ARTIFICIAL INTELLIGENCE: 7 QUESTIONS (AND ANSWERS)

WHAT IS ARTIFICIAL INTELLIGENCE?

Artificial intelligence (AI) refers to the leveraging of multiple technologies that together create a device or construct that accomplishes certain tasks formerly requiring human input. In higher education, the principles of AI underlie a range of innovative systems, including analytics, robot writers, virtual experiences, and intelligent tutoring systems.

1



WHAT ARE THE IMPLICATIONS FOR TEACHING AND LEARNING?

AI bots can respond to student questions when access to the instructor and teaching assistants is limited or unavailable. AI has the potential to give every student a computer-simulated personal mentor and provide better communication between classrooms worldwide by offering translation services and cultural context. Lectures may be accompanied or augmented by immersive virtual reality environments populated by AI personalities that offer safe opportunities to practice emerging skills.

7



HOW DOES IT WORK?

To exhibit intelligence, computers apply algorithms to find patterns in large amounts of data—a process called machine learning, which plays a key role in a number of AI applications. AI systems often incorporate human feedback to help calibrate the system's learning.

2



WHERE IS IT GOING?

AI will trend toward devices and constructs that conform more closely to human behavior and systems that are better able to handle conflicting or false information. The use of AI systems and devices will expand further into routine activities. Users may come to understand AI as a system that enhances human capabilities in a partnership between humans and machines, leveraging what each does best.

6



WHO'S DOING IT?

Many colleges and universities are developing AI projects that aid teaching and learning, such as the Pennsylvania State University, Georgia Tech, MIT, and Harvard.

3



WHAT ARE THE DOWNSIDES?

Considerable misunderstanding exists about what AI can and cannot do, resulting in inflated expectations and a risk that users could assign inappropriate kinds and amounts of authority to AI systems. For AI developers, one key issue is an emerging lack of transparency among corporate entities that see their AI programming and algorithmic development as intellectual property.

5



WHY IS IT SIGNIFICANT?

AI opens the possibility of individual tutoring to students who could never otherwise have access to it. AI learning agents have the potential to function like adaptive learning but at a much more sophisticated and nuanced level. AI allows faculty and students to do their work more effectively by providing not just tutors but AI assistants for scheduling, interactive immersive simulations, and human-machine partnerships.


4



This infographic was adapted from *7 Things You Should Know About Artificial Intelligence in Teaching and Learning* (Louisville, CO: EDUCAUSE Learning Initiative, April 2017).
© 2017 EDUCAUSE. The text of this work is licensed under a [Creative Commons BY-NC-ND 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/).

ARTIFICIAL INTELLIGENCE: THREAT OR OPPORTUNITY?

BY BRIAN FLEMING



COLIN ANDERSON/ISTOCKSY

T

he last few years have been rough for higher education. According to a 2018 Gallup Poll that tracked Americans' confidence in colleges and universities, over the previous five years higher education saw its sharpest decline in public trust, with only 48 percent of those surveyed expressing confidence, down from 57 percent in 2015.¹

But statistics like these can be overstated. Americans distrust many traditional institutions these days: not only higher education but also government and the media. That distrust extends to big technology companies such as Facebook and Google. According to the Edelman Trust Barometer 2020, which tracks consumer sentiment across a range of sectors, Americans distrust—or are at least ambivalent about—the development of advanced technologies such as artificial intelligence (AI) by companies that may not be positively and responsibly shaping our future.²

Think about the fallout from Facebook's Cambridge Analytica debacle, in which millions of users' profiles were harvested without consent and used for political advertising. And consider Uber's Advanced Technologies Group, which had no official safety plans in place when one of its self-driving test cars crashed and killed a woman. These examples are frightening because they

appear to be void of responsible leadership acting in the public's collective best interests. They leave us not knowing who we can trust in a brave new world. There is, however, one exception—according to a 2019 survey from the University of Oxford's Future of Humanity Institute (FHI), which asked 2,000 Americans to rate their confidence in actors developing artificial intelligence. Half of Americans surveyed said they trusted higher education (and the military) above all (more than government

agencies, non-profit research collaboratives, and big technology companies) to build, manage, and govern artificial intelligence.³

We should lean into this finding. It not only signifies at least a pocket of trust remaining in higher education institutions but also offers an opportunity for college/university researchers, faculty, staff, and administrators to regain lost ground and exemplify AI leadership at a time when our institutions—

and our world—need us most.

Leadership is increasingly digital in focus and is present in just about every sector today. Generally, digital leadership describes an emerging class of roles, responsibilities, and competencies needed to lead organizations in a digital world. But we should not confuse digital leaders with digital evangelists, at least not in higher education.

Digital leaders are equipped to lead in a digital world. They understand its


We need leaders who are engaged in the world of artificial intelligence—whether as researchers, subject-matter experts, educators, ethicists, or administrators in our communities and the world at large—and who are committed to building transparency and trust within the AI world.

complexity and also the dissonance and distrust that digital can create, and they help others make meaning within and out of it. Good digital leaders are virtuous and altruistic. According to Deborah Ancona, who studies digital leadership at the MIT Sloan School of Management, digital leaders are sense-makers who help others “create meaning out of the messy world.”⁴ Their lens is digital, but their focus is human.

We need more digital leaders in higher education who are sense-makers not only for their own institutions but for the public at large. We need leaders who are optimistic about this technology but also cautious. We need leaders who are engaged in the world of artificial intelligence—whether as researchers, subject-matter experts, educators, ethicists, or administrators in our communities and the world at large—and who are committed to building transparency and trust within the AI world.

This is something technology companies struggle to do, but it's in the DNA of higher education. Think of digital leadership as a strategy of engagement, taking the understanding of, resources for, and experiences with artificial intelligence cultivated within colleges and universities—whether through basic research, experimentation, teaching, or academic innovation—out into the world to meet its most pressing challenges. Doing so not only will quell fears but also may instill—perhaps even increase—confidence in higher education at a time when we need it most.

As artificial intelligence continues to move further into the mainstream (which it will) and as regulators struggle to govern AI research and development (which they will) and as the market continues to coalesce around big-tech companies such as Facebook and Google (which it will), higher



COLIN ANDERSON/ISTOCK

education is uniquely poised to gain public trust once again. ■


Notes

1. Jeffrey M. Jones, “[Confidence in Higher Education Down Since 2015](#),” Gallup (website), October 9, 2018.
2. Edelman Trust Barometer 2020, “[Special Report: Trust in Technology](#).”
3. Baobao Zhang and Allan Dafoe, “[Artificial Intelligence: American Attitudes and Trends](#)” (Center for the Governance of AI, Future of Humanity Institute, University of Oxford, January 2019).
4. Deborah Ancona, “[Five Rules for Leading in a Digital World](#),” *MIT Sloan Management Review*, October 28, 2019.

Brian Fleming is Associate Vice Chancellor, Learning Ecosystem Development, at Northeastern University. At the time of publication, he was Vice President for Innovation and Strategy at Southern New Hampshire University (SNHU).

© 2020 Brian Fleming.

This article was originally published in *EDUCAUSE Review* on [May 18, 2020](#).



ARTIFICIAL INTELLIGENCE

AND ETHICAL ACCOUNTABILITY

BY LINDA FISHER THORNTON

W

hile artificial intelligence (AI) is a potential game-changer for improving efficiency, it may also come with ethical baggage. John O'Brien, president and CEO of EDUCAUSE,

defines digital ethics as “doing the right thing at the intersection of technology innovation and accepted social values.”¹ IT departments that want to tap into the benefits of AI can't assume that AI-enabled programs and services on the market will live up to that definition. The vendor that created the product or service may not have considered ethics during each stage of development. Or, perhaps the vendor adhered to high ethical standards but is bought out by a company that is guided by a different set of moral principles. IT departments should be ready to assess and manage ethics before, during, and after AI deployment.

Preparing and Taking Responsibility for AI Ethics

Universities are already using AI, possibly without having considered the risks fully. According to the *2020 EDUCAUSE Horizon Report Teaching and Learning Edition*, AI is broadly in use across higher education campuses, embedded in “test generators, plagiarism-detection systems, accessibility products, and even common word processors and presentation products.”²

AI can provide game-changing improvements by automating processes that are typically handled by humans, but the areas of risks associated with AI are broad and deep. Potential risks include “everything from AI algorithmic bias and data privacy issues to public safety concerns from autonomous machines running on AI.”³ Changes that occur after deployment and are outside of the IT department's

control may also increase risk. For example, a university's learning management system (LMS) may start with a robust data privacy policy, but the company that manages it may be sold to a private company, resulting in unforeseen risks to student data.⁴

Ensuring Ethical AI

Determining who is responsible for ethical AI turns out to be more complicated than identifying the person who created the program. There are potentially multiple responsible parties, including programmers, sellers, and implementers of AI-enabled products and services. For AI to be ethical, multiple parties must fulfill their ethical obligations.

The AI Programmer and the Programmer's Manager

Lynn Richmond, an associate at BTO Solicitors in Edinburgh, Scotland, writes, “Where AI is employed to mimic the process carried out by a human . . . the party who has provided the wrong data, the wrong pre-determined result, or the wrong process is likely to be liable.”⁵ It seems logical that the programmer would be assumed responsible if anyone is harmed by AI. The programmer, however, was not acting alone. The programmer's manager and others who helped create the AI are also accountable.

John Kingston, a senior lecturer in cybersecurity at Nottingham Trent University in England, notes that determining accountability may include “debates [about] whether the fault lies with the programmer; the program designer; the expert who provided the knowledge; or the manager who appointed the inadequate expert, program designer, or programmer.”⁶ With this high level of accountability for developing ethical AI, some vendors are looking into how to track the thought

AI can provide game-changing improvements by automating processes that are typically handled by humans, but the areas of risks associated with AI are broad and deep.

process that the AI is using so they can better prepare in case a legal defense becomes necessary.⁷

The AI Marketer and Salesperson

Once an AI program is on the market, whose job is it to ensure that it will have the intended outcome and will not just automate ethical missteps? Chris Temple, a partner in the litigation practice at Fox Rothschild, names “sellers” on the list of the many parties who could be seen as legally responsible for nonhuman decision-making if things go wrong.⁸

But just because sellers have a legal responsibility doesn't necessarily mean that the programs they are marketing and selling are ethical. Chrissy Kidd, a freelance technology writer for BMC Software, reminds us that “some consider ‘as a service’ offerings a black box—you know the input and the output, but you don't understand the inner-workings.”⁹ This lack of understanding creates a major issue for the end user because IT departments shouldn't simply trust that AI programs on the market will already have ethics “built into” their design and sales processes.

Higher Education Institutions That Implement AI

The responsibility for ethical AI is shared by many different stakeholders, from programmers and their managers to the companies that market and sell the programs—and even the higher education institutions that deploy AI products and services.

Whether or not ethical design was used when the AI was developed (or whether ethical practices were followed when it was sold) can become a problem for the end user. For example, new legislation in Illinois places some of the ethical responsibility associated with AI video hiring with the company purchasing the tool.¹⁰ IT departments will need to determine whether or not ethics were carefully applied in the planning, programming, and selling of an AI-enabled program before purchasing the program.

What can IT departments do to be ready to take on this level of ethical responsibility? Here are five practical steps IT departments can take to manage the ethics of AI. These steps should be applied to AI that is





EDUCAUSE

ANNUAL CONFERENCE 2022

October 25–28
Denver, CO

November 2–3
Online

Join us in Denver, CO or online.

events.educause.edu/annual-conference

already deployed on campuses as well as AI that is being considered for future use.

Five Steps IT Departments Can Take to Manage the Ethics of AI

- 1. Make AI ethics a priority.**
 - Commit to understanding the ethical issues of AI deployment.
 - Consider ethical issues when making decisions so that AI use is more likely to lead to ethical outcomes.
- 2. Consider current and future AI applications across higher education functions.**
 - Ask, “Where is AI already in use on our campuses, and where might it be used in the future?”
 - Read about important overarching [ethical principles](#) and [guidelines](#).
 - Consult sources with specific guidance on [ethical design](#).
- 3. Create a custom AI ethics code to guide decision-making.**
 - Review these articles: [IBM](#), [Sage](#), [Intel](#), [World Economic Forum](#), [Deloitte](#), and [Nesta](#).
 - Review this AI ethics code example: [Society of Corporate Compliance and Ethics](#).
- 4. Apply the AI ethics code to existing AI deployments.**
 - Use the AI ethics code to evaluate and manage the ethical risks of AI that is already in use.
- 5. Apply the AI ethics code to future AI deployments.**
 - Use the AI ethics code to select AI-enhanced programs and plan future AI deployments to avoid bias and [unintended consequences](#).

Taking responsibility for AI means considering the ethical issues at every

step. This kind of ethical thinking must be applied methodically whether designing, marketing, selling, or deploying AI solutions, which means that when selecting AI-enabled programs for hiring, teaching, testing, or student support, institutions must intentionally weigh the cost savings and efficiency the programs bring against the risks of human harm in areas including fairness, privacy, bias, and public safety. ■

Notes

1. John O'Brien, [“Digital Ethics in Higher Education: 2020,”](#) *EDUCAUSE Review*, May 18, 2020.
2. Malcolm Brown, et al., [2020 EDUCAUSE Horizon Report, Teaching and Learning Edition](#), research report, (Louisville, CO: EDUCAUSE, 2020).
3. John McClurg, [“AI Ethics Guidelines Every CIO Should Read,”](#) *Information Week*, August 7, 2019.
4. Jeffrey Young, [“As Instructure Changes Ownership, Academics Worry Whether Student Data Will Be Protected.”](#) *EdSurge*, January 17, 2020.
5. Lynn Richmond, [“Artificial Intelligence: Who’s to Blame?”](#) SCL, August 8, 2018.
6. John Kingston, [“Artificial Intelligence and Legal Liability,”](#) in *Research and Development Systems XXXIII* (New York: Springer, 2016).
7. Joseph Lazzarotti, [“Illinois Leads the Way on AI Regulation in the Workplace,”](#) *National Law Review*, October 24, 2019.
8. Chris Temple, [“AI-Driven Decision-Making May Expose Organizations to Significant Liability Risk,”](#) *Corporate Compliance Insights*, September 11, 2019.
9. Chrissy Kidd, [“AI as a Service \(AlaaS\): An Introduction,”](#) *Machine Learning & Big Data Blog* (blog), BMC Blogs, April 25, 2018.
10. John Edwards, [“AI Regulation: Has the Time Arrived?”](#) *Information Week*, February 24, 2020.

Linda Fisher Thornton is CEO of Leading in Context LLC and an Adjunct Associate Professor at the University of Richmond.

© 2020 Linda Fisher Thornton. The text of this work is licensed under a [Creative Commons BY-NC-ND 4.0 International License](#).

This article was originally published in *EDUCAUSE Review* on [July 31, 2020](#).



LIGHTSPRING/SHUTTERSTOCK

EDUCAUSE QuickPoll Results

Artificial Intelligence Use in Higher Education

By D. Christopher Brooks

EDUCAUSE is helping institutional leaders, IT professionals, and other staff address their pressing challenges by gathering and sharing data. This report is based on an EDUCAUSE QuickPoll. QuickPolls enable us to rapidly gather, analyze, and share input from our community about specific emerging topics.¹

The Challenge

We are on the verge of peak hype about how artificial intelligence (AI) can (and will) transform our lives. No fewer than seven emerging AI technologies were prominently featured on “The Gartner Hype Cycle for Emerging Technologies, 2020.”² Several technologies on EDUCAUSE’s “The Top 10 Strategic Technologies for 2020” explicitly incorporate or are reliant on AI.³ And while AI might seem to be a technology in search of a campus, some promising applications have been emerging in domains such as teaching and learning, student success, and accessibility.⁴



But how widespread is the use of AI in higher education today? In this QuickPoll, we operationalized Elana Zeide’s categories of AI applications in higher education⁵ to better understand how and how widely AI is being used for institutional tasks, student success and support tasks, and instructional tasks.⁶

The Bottom Line

AI is most developed for instructional use, especially for monitoring student behavior during exams and ferreting out plagiarism. AI is being used the least for institutional tasks. Significant numbers of respondents reported that they don’t know the status of AI at their institutions across all categories, suggesting that AI use may be

obscure and/or intangible.⁷ Immature data governance, concerns about algorithmic bias, and ineffective data management and integration pose the greatest challenges to the implementation of AI in higher education. For now, the hype surrounding the revolutionary impact of AI on higher education appears to be just that—hype.

The Data: Instructional Use

AI is being used to monitor students and their work. The most prominent uses of AI in higher education are attached to applications designed to protect or preserve academic integrity through the use of plagiarism-detection software (60%) and proctoring applications (42%) (see figure 1). Although both applications, especially the former, have been in use for some time, the latter has experienced considerable growth due to the expansion of online learning during the pandemic. Both types of tools have come under scrutiny for violating the privacy of students and producing false positives; the use of proctoring software is also associated with a litany of problems related to exam performance due to anxiety, technology failures, and socioeconomic and racial bias.⁸ Responding to some of these concerns, one provider recently announced that it will no longer provide systems based solely on AI, requiring a human being to analyze the captured video.⁹

AI is not going to replace instructors anytime soon. Most respondents reported that AI is not in use at their institutions as it relates to instructional tasks, excepting plagiarism-detection software and proctoring applications. Majorities of respondents told us that AI is not in use—and that there are no plans to use it in the future—for key instructional tasks such as providing feedback on assignments, tutoring, conducting assessments, and grading assignments. Although substantial percentages of respondents told us that their institution is tracking AI for these tasks, usage appears to be limited.



LIGHTSPRING/SHUTTERSTOCK

Figure 1. AI Usage for Instructional Tasks

	Not in use, no plan to use	Tracking for potential use	Planning, piloting, and initial usage	Some usage	Full or mostly full usage
Providing feedback on assignments	58%	21	7	14	0
Tutoring	54	33	3	7	2
Conducting assessments	53	26	6	15	0
Grading assignments	52	22	4	22	1
Creating personalized learning opportunities	44	35	5	15	1
Proctoring	37	16	5	35	6
Using plagiarism-detection software	22	10	9	39	20

The Data: Student Success and Support Use

The chatbots are coming! The chatbots are coming! A sizable percentage (36%) of respondents reported that chatbots and digital assistants are in use at least somewhat on their campuses, with another 17% reporting that their institutions are in the planning, piloting, and initial stages of use (see figure 2). The use of chatbots in higher education by admissions, student affairs, career services, and other student success and support units is not entirely new, but the pandemic has likely contributed to an increase in their use as they help students get efficient, relevant, and correct answers to their questions without long waits.¹⁰ Chatbots may also liberate staff from repeatedly responding to the same questions and reduce errors by deploying updates immediately and universally.

Student success tools are a potential area of growth for AI. A limited but comparatively sizable group of respondents reported that AI is being used for student success tools such as identifying students who are at-risk academically (22%) and sending early academic warnings (16%); another 14% reported that their institutions are in the stage of planning, piloting, and

initial usage of AI for these tasks. That said, these numbers seem low, given that student success tools have been around for nearly a decade and are deployed widely.¹¹ One possible explanation for this discrepancy is semantic—some might not view the analytics that power many student success tools as AI when, in fact, analytics is a type or subset of AI.¹²

The Data: Institutional Use

AI is sparsely used for institutional tasks. Most respondents reported that AI is not in use at their institutions as it relates to institutional tasks (see figure 3). Clear majorities of respondents reported a complete lack of interest in using AI for institutional tasks such as planning curricula, making or contributing to financial aid decisions, development and fundraising, and making or contributing to admissions decisions. The tasks with the most AI use are nudging accepted applicants to put down deposits (17%), planning academic support resources (15%), and marketing and recruiting (15%). Some respondents identified additional tasks that use AI, including the analysis of student evaluations of teaching, instructional planning, social media analysis, support desk services, and attendance on campus.

Figure 2. AI Usage for Student Success and Support Tasks

	Not in use, no plan to use	Tracking for potential use	Planning, piloting, and initial usage	Some usage	Full or mostly full usage
Assessing student engagement by facial expressions	82%	13	2	1	1
Providing emergency financial aid	67	26	2	2	2
Assessing financial need	61	28	5	4	3
Recommending classes	49	34	8	8	1
Identifying students who are struggling with nonacademic issues	44	33	12	9	3
Answering questions about academic progress	43	36	11	10	1
Using adaptive learning materials	40	34	14	10	1
Sending early academic warnings	32	38	14	14	2
Identifying students who are at-risk academically	30	33	14	19	3
Using chatbots and digital assistants	21	26	17	25	11

The Data: What You Don't Know... Can what you don't know hurt you?

Significant percentages of respondents reported that they don't know the status of AI at their institutions across all categories. Ranges of "don't know" responses:

- **Instructional tasks:** 8% (using plagiarism-detection software) to 23% (tutoring)
- **Institutional tasks:** 20% (planning academic support resources) to 33% (development and fundraising)
- **Student success and support tasks:** 10% (using chatbots and digital assistants) to 32% (assessing financial need)

The lack of knowledge about AI on one's campus could be attributed to a vague or incorrect understanding of what AI is, an inability to observe AI work (because it tends to be baked into applications and tools), a lack of awareness of the ways in which AI might be used in different units across campus, and/or an actual lack of AI usage on campus. Regardless, that such large

percentages responded with "don't know" suggests that the importance of AI to higher education may be presently overstated.

Common Challenges

We're just not ready. About two-thirds of respondents reported that institutional deficiencies to support the adoption and maintenance of AI are the main challenges to the implementation of AI at their institutions (see figure 4). Nearly three-quarters of respondents said that ineffective data management and integration (72%) and insufficient technical expertise (71%) present at least a moderate challenge to AI implementation. Financial concerns (67%) and immature data governance (66%) also pose challenges. Insufficient leadership support (56%) is a foundational challenge that is related to each of the previous listed challenges in this group.

Show me the ethics! Concerns about ethics related to AI use (68%) and concerns about algorithmic bias (67%) pose significant challenges to AI implementation. Echoing the findings of Safiya Umoja

Figure 3. AI Usage for Institutional Tasks

	Not in use, no plan to use	Tracking for potential use	Planning, piloting, and initial usage	Some usage	Full or mostly full usage
Planning curricula	66%	22	5	6	1
Making or contributing to financial aid decisions	60	27	6	7	1
Development and fundraising	57	27	7	8	1
Making or contributing to admissions decisions	53	28	10	8	1
Nudging accepted applicants to put down deposits	47	26	10	13	4
Planning academic support resources	45	28	12	12	3
Marketing and recruiting	42	30	13	14	1

Noble’s book *Algorithms of Oppression*, one respondent whose campus primarily serves minority populations told us that, “Bias issues in AI are rampant. As it stands now, [using AI] would have too great a negative impact on our students.” Another expressed concerns that “AI has too much bias built in that is very difficult to remove or mitigate.” Risk to institutional reputation poses a challenge as well, but what remains unclear is whether respondents see having AI implemented on campus is desirable . . . or derisible. Figuring out how, if at all, AI aligns with current institutional missions is the least threatening concern.

Promising Practices

Current use of AI is a mile wide and an inch deep. We asked respondents to share some promising practices in the use of AI at their institutions. The responses run the gamut of tasks identified above and a few that we hadn’t considered:

- Chatbots for informational and technical support, HR benefits questions, parking questions, service desk questions, and student tutoring
- Research applications, conducting systematic reviews and meta-analyses, and data science research

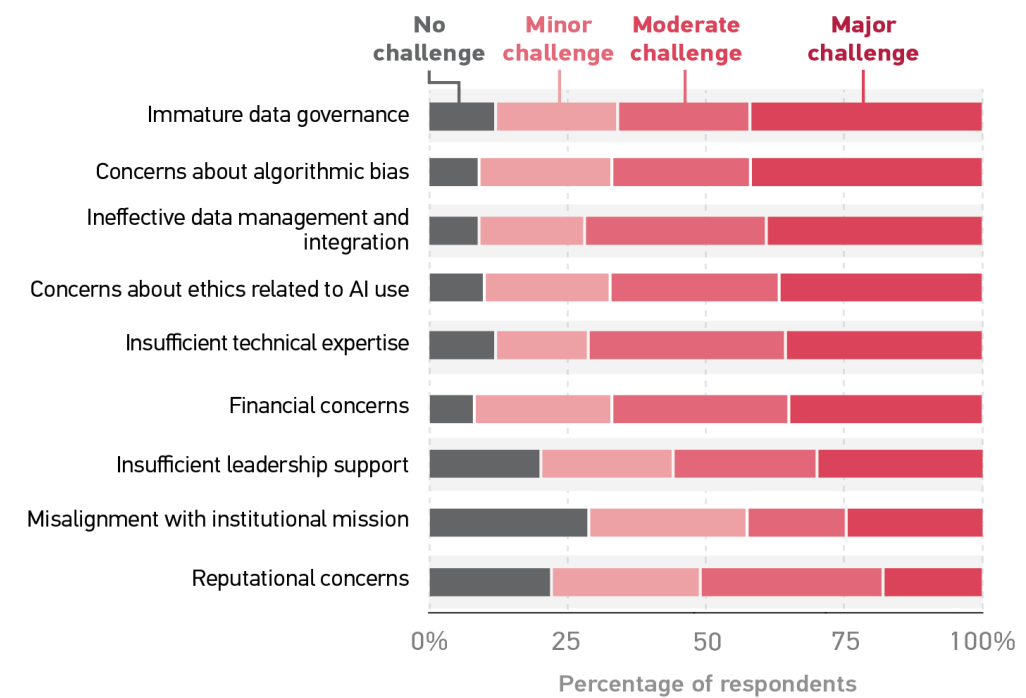
- Library services
- Recruitment of prospective students
- Providing individual instructional material pathways, assessment feedback, and adaptive learning software
- Proctoring and plagiarism detection
- Student engagement support and nudging, monitoring well-being, and predicting likelihood of disengaging the institution
- Detection of network attacks
- Recommender systems

All QuickPoll results can be found on the [EDUCAUSE QuickPolls web page](#). For more information and analysis about higher education IT research and data, please visit the [EDUCAUSE Research web page](#).

Notes

1. QuickPolls gather data in a single day instead of over several weeks, are distributed by EDUCAUSE staff to relevant [EDUCAUSE Community Groups](#) rather than via our enterprise survey infrastructure, and do not enable us to associate responses with specific institutions.
2. Kasey Panetta, “5 Trends Drive the Gartner Hype Cycle for Emerging Technologies, 2020,” Gartner, August 18, 2020.
3. Mark McCormack, D. Christopher Brooks, and Ben Shulman, *Higher Education’s 2020 Trend Watch and Top 10 Strategic Technologies*, research report (Louisville, CO: ECAR, January 2020). See [The Top 10 Strategic Technologies for 2020](#).

Figure 4. Common Challenges to the Implementation of AI



4. Bryan Alexander, “5 Ais in Search of a Campus,” *EDUCAUSE Review*, October 14, 2019; Kathe Pelletier, Malcolm Brown, D. Christopher Brooks, Mark McCormack, Jamie Reeves, and Nichole Arbino, with Aras Bozkurt, Steven Crawford, Laura Czerniewicz, Rob Gibson, Katie Linder, Jon Mason, and Victoria Mondelli, *2021 EDUCAUSE Horizon Report, Teaching and Learning Edition* (Boulder, CO: EDUCAUSE, 2021); Thomas Miller and Melissa Irvin, “Using Artificial Intelligence with Human Intelligence for Student Success,” *EDUCAUSE Review*, December 9, 2019; and Judy Brewer, Carly Gerard, and Mark Hakkinen, “The Impact of AI on Accessibility,” *EDUCAUSE Exchange*, *EDUCAUSE Review*, November 4, 2020.
5. Elena Zeide, “Artificial Intelligence in Higher Education: Applications, Promise and Perils, and Ethical Questions,” *EDUCAUSE Review*, August 26, 2019.
6. The poll was conducted on June 7–8, 2021, consisted of eight questions, and resulted in 195 responses. Poll invitations were sent to participants in EDUCAUSE community groups focused on IT leadership. Our sample represents a range of institution types and FTE sizes, and most respondents (88%) represented U.S. institutions.
7. For this report, cited percentages are among those respondents who reported knowing the status of AI at their respective institution.
8. Daniel Woldeab and Thomas Brothen, “21st Century Assessment: Online Proctoring, Test Anxiety, and Student Performance,” *International Journal of E-Learning & Distance Education* 34, no. 1 (2019); D. Christopher Brooks, *Student Experiences Learning with Technology in the Pandemic*, research report (Boulder, CO: EDUCAUSE, April 2021); Shea Swauger, “Software That Monitors Students During Tests Perpetuates Inequality and Violates Their Privacy,” *MIT Technology Review*, August 7, 2020; and Todd Feathers, “Proctorio Is Using Racist Algorithms to Detect Faces,” *Vice*, April 8, 2021.
9. Scott Jaschik, “ProctorU Abandons Business Based Solely on AI,” *InsideHigherEd*, May 24, 2021.
10. D. Christopher Brooks, “EDUCAUSE QuickPoll Results: Student Success Technologies,” *EDUCAUSE Review*, April 9, 2021.
11. See “The EDUCAUSE Student Success Almanac.”
12. Zeide, “Artificial Intelligence in Higher Education.”

D. Christopher Brooks is Senior Analyst at the Tambellini Group. At the time of publication, he was Director of Research at EDUCAUSE.

© 2021 D. Christopher Brooks. The text of this work is licensed under a [Creative Commons BY-NC-ND 4.0 International License](#).

This article was originally published in *EDUCAUSE Review* on [June 11, 2021](#).

How Higher Ed CIOs Can Leverage AI

By Toby Jackson

SPONSORED CONTENT



In higher education, as in the business world, IT organizations are often viewed as the worker bees. The IT staff buzz around campus tinkering with the Wi-Fi, fixing the overhead projectors in the classrooms, and taking orders from the queen bees in facilities, finance, and the provost's office.

But just as chief information officers in the corporate world are now being pushed to leverage technology for a variety of strategic purposes—from making products more efficiently to finding new customers—CIOs in higher education are also facing digital transformation. This change reflects the mounting pressure facing institutional leaders due to a decade-long enrollment decline that accelerated during the COVID-19 pandemic.¹ Leaders at public and private colleges and universities—squeezed by state legislatures, debt-wary students, and now inflation—are worried about their annual budgets and long-term viability of their institutions. Increasing student success metrics, bolstering the mental health of students and employees, and improving the effectiveness of diversity and equity initiatives remain ongoing priorities.²

Technology can play a critical role in helping institutions navigate an increasingly uncertain future, but for that to happen, higher ed CIOs must take a more active role in guiding the digital transformation happening on their campuses. They'll need all the tools they can muster, and CIOs who want to be change agents on their campus should make sure they have artificial intelligence in their toolkits.


Consider what we've already seen in the corporate world about the potential impact of AI in action. In investment finance, for instance, a top-notch analyst can sift through tons of data, pick the most important numbers to focus on,

and suss out trends. But great analysts are hard to find. And this process—gathering data and then figuring out the correct data—is extremely hard to scale across multiple exchanges in stocks, commodities, and currencies. Smart companies have deployed machine learning models that work alongside humans to point out patterns in data: things that happened in the past and that now might be taking place again. In this way, AI finds a potential needle in the proverbial haystack—one that analysts can then pull out for in-depth review.

Higher education—which, at its core, is all about the pursuit of knowledge—has the potential to also use AI in this digital evolution. Colleges and universities are awash in data. All of these institutions have learning management systems, student information systems, customer relationship management systems for admissions, and alumni affairs and systems that manage finances, personnel, security, and a host of other crucial campus functions. Each of those systems is a data repository. Every time a prospective student, a current student, or a graduate fills out an online form, the institution adds more data to one of its systems.

But data by itself doesn't tell us very much. It needs more context to become information. And it needs to be connected to other data to become knowledge. Then it needs to be understood so that campus leaders can use it to make smart decisions that will, for instance, increase enrollment, retain more students from year to year, and help students stay on track to graduate on time.

Presidents and deans and administrators have long relied on their own expertise as they have made decisions that affect admissions, academics, and graduation rates. But as colleges and universities get more complex and confront more



DEN RISE /SHUTTERSTOCK

complicated issues amid increased demands for better outcomes, they need high-tech assistance in the form of AI to parse and connect and understand all of this data they've collected.

So how can campus CIOs leverage AI to be more proactive and propose solutions to these challenges?

First: CIOs can become familiar with how other businesses and industries have applied AI and draw inspiration from that. In marketing, for instance, AI helps identify potential customers. On campus, AI can help bolster the admissions office by finding prospective students who might be interested in applying and matriculating.

Second: CIOs can ask what their institution is trying to achieve. What are the challenges? Where are the pain points? CIOs should hunt for the problems on campus and try to explore where AI can play a part in helping to solve them.

Third: CIOs can use AI not only to deal with challenges that exist today but also to anticipate issues that might arise tomorrow and figure out ways to address them. In this way, CIOs can be problem-solvers, not just solution implementers.

This digital evolution from data to information to knowledge—assembling the pieces, connecting them, and getting a broader view of the world—is what enables action and drives outcomes. AI can play a significant role in making this transformation possible in higher education. ■

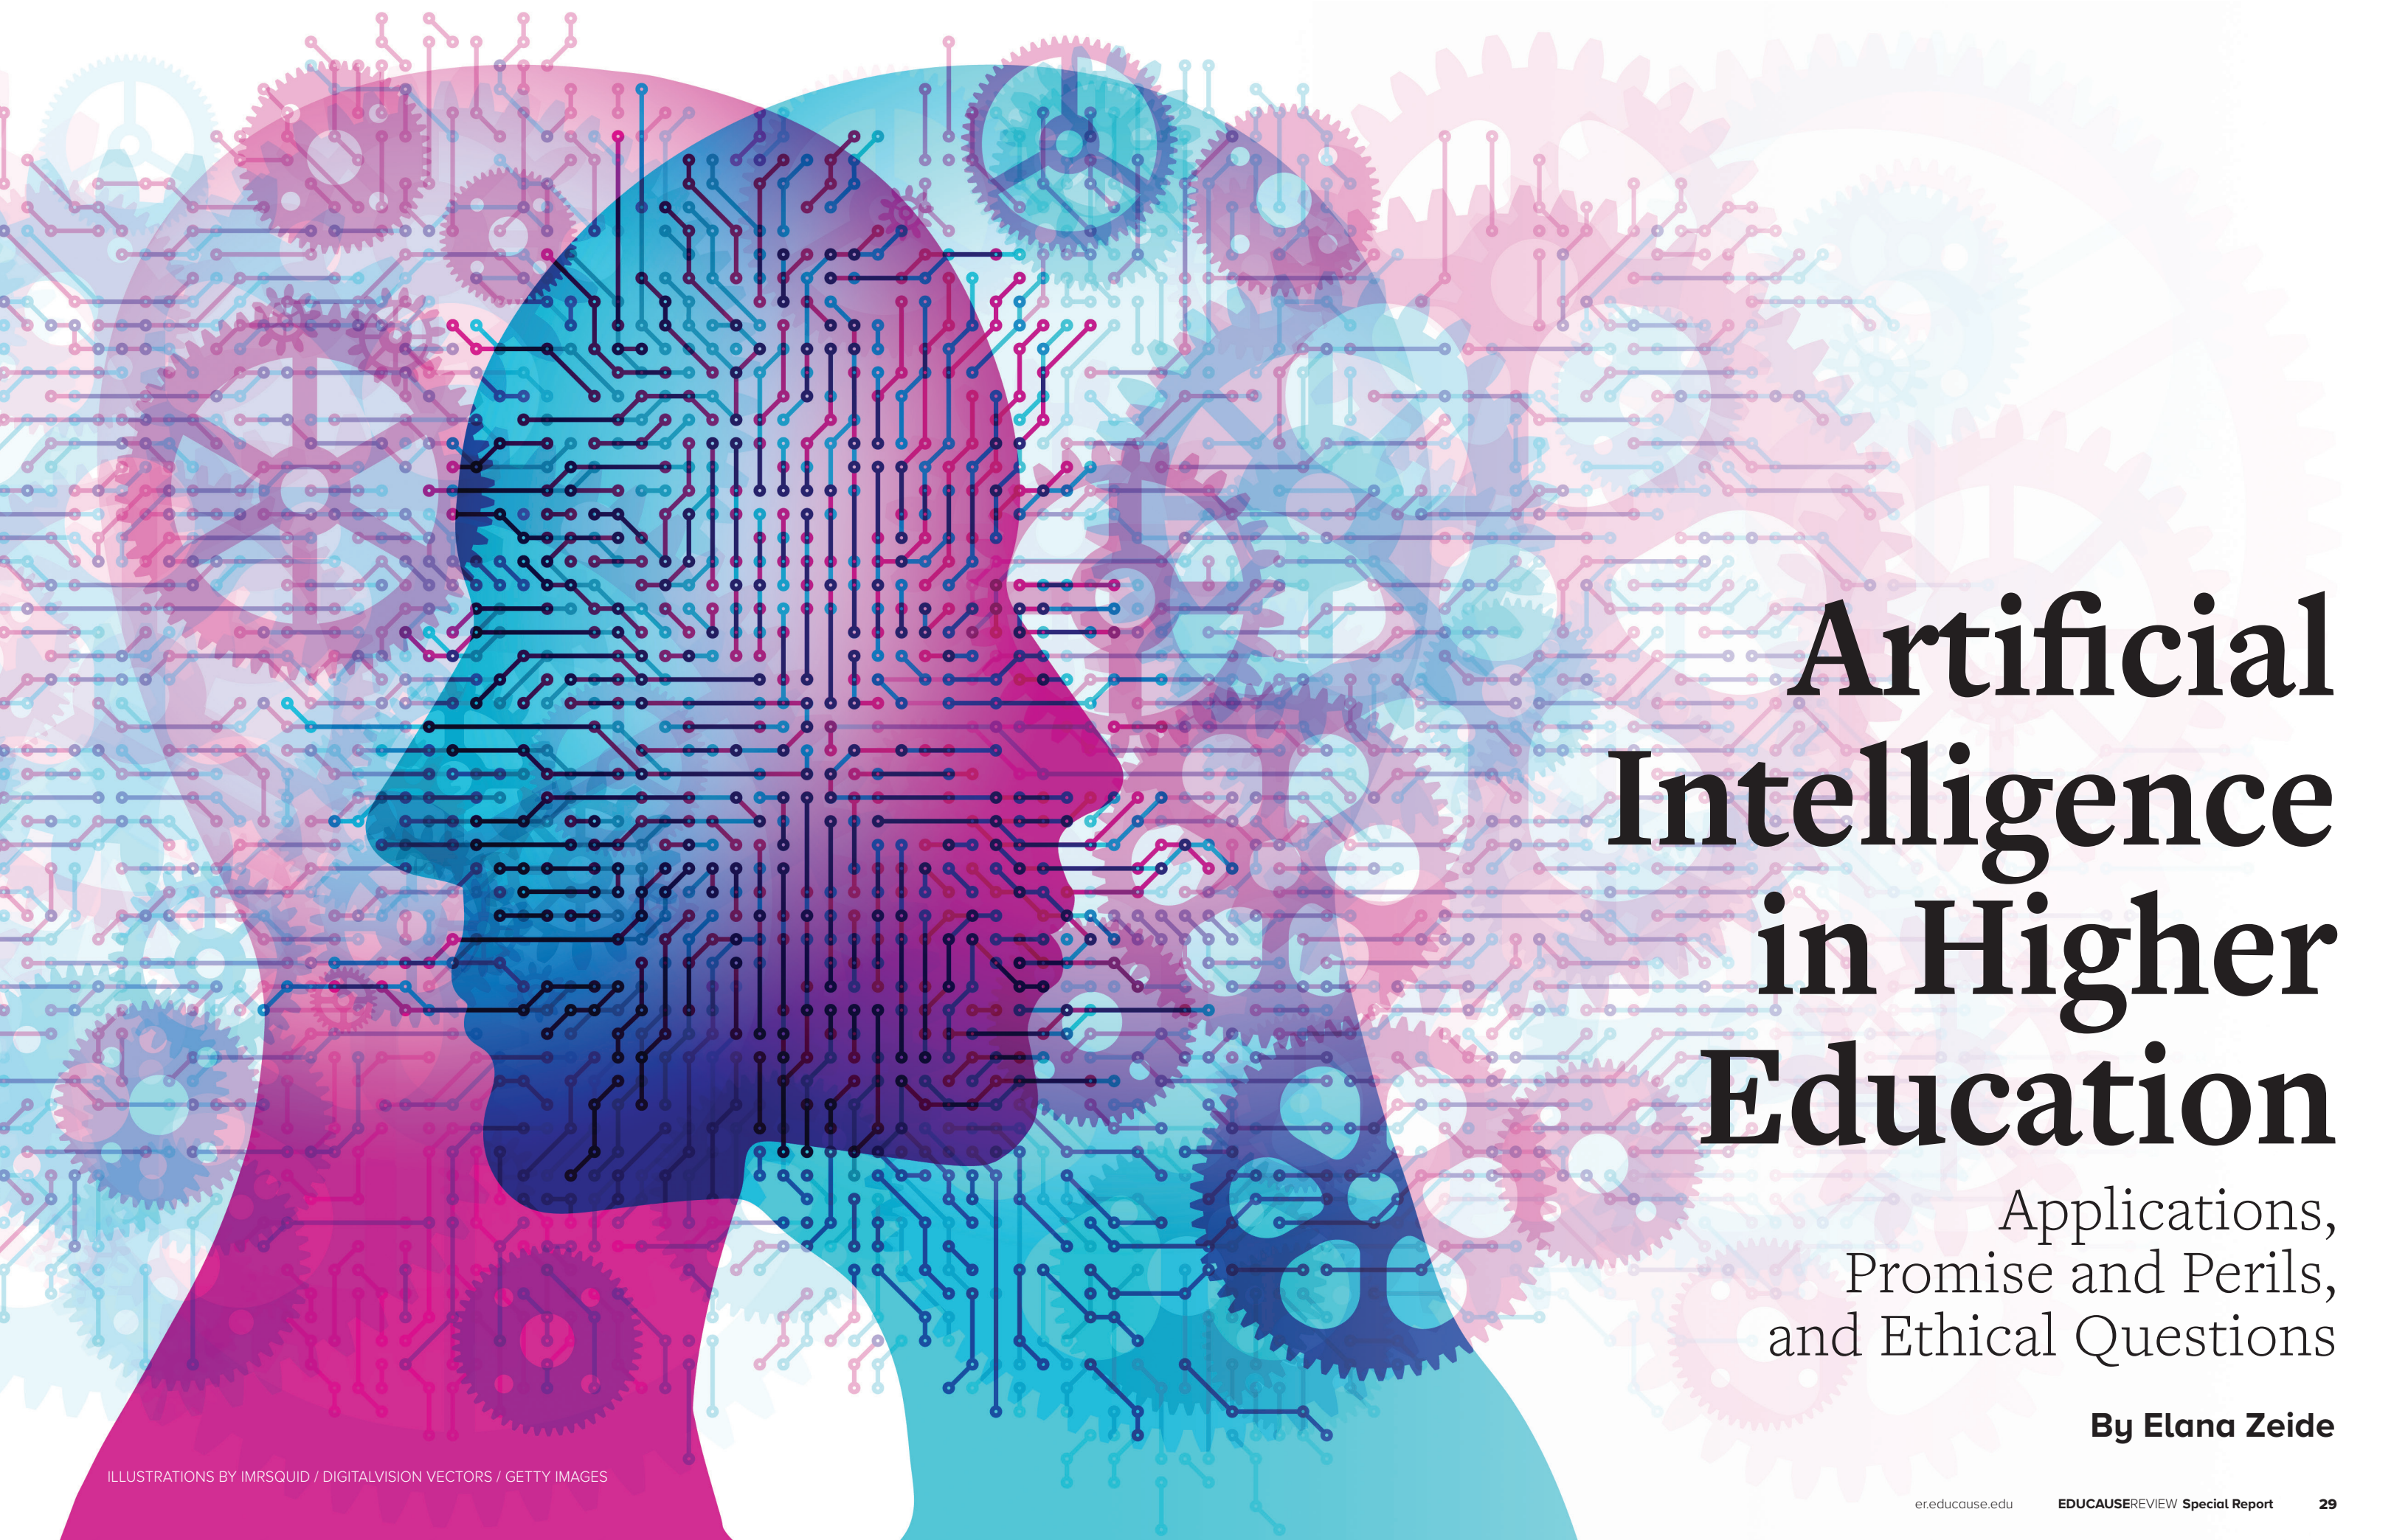
Notes

1. Jon Marcus, "[Drop in College Enrollment Threatens to Cause Long-Term Economic, Social Consequences](#)," *Washington Post*, January 22, 2022.
2. "[Presidential Priorities: Areas of Concern for College and University Presidents](#)," *Higher Ed Dive*, June 10, 2020; Michael Vasquez, "[As the Pandemic Grinds On, Here Are 5 Big Worries of College Presidents](#)," *Chronicle of Higher Education*, October 12, 2020; "[Campus Priority: Advancing Diversity, Equity, and Inclusion](#)," American Association of Colleges and Universities (website), accessed April 25, 2022.

Toby Jackson is Chief Technology Officer at Mainstay.

© 2022 Mainstay

This article was originally published in *EDUCAUSE Review* on [June 1, 2022](#).



Artificial Intelligence in Higher Education

Applications,
Promise and Perils,
and Ethical Questions

By Elana Zeide

What is artificial intelligence?

In any discussion of artificial intelligence (AI), this is almost always the first question. The subject is highly debated, and I won't go into the deep technical issues here. But I'm also starting with this question because the numerous myths and misconceptions about what artificial intelligence is, and how it works, make considering its use seem overly complex.

When people think about artificial intelligence, what often comes to mind is The Terminator movies. But today we are far from machines that have the ability to perform the myriad of tasks even babies shift between with ease—although how far away is a matter of considerable debate. Today's artificial intelligence isn't general, but narrow. It is task-specific. Consider the computer program that famously beat the world's champion in the Chinese game Go. It would be completely befuddled if someone added an extra row to the playing board. Changing a single pixel can throw off image-recognition systems.

Broadly, *artificial intelligence* is the attempt to create machines that can do things previously possible only through human cognition. Computer scientists have tried many different mechanisms over the years. In the last wave of AI enthusiasm, technologists tried emulate human knowledge by programming extensive rules

into computers, a technique called expert systems. Today's artificial intelligence is based on machine learning. It is about finding patterns in seas of data—correlations that would not be immediately intuitive or comprehensible to humans—and then using those patterns to make decisions. With “predictive analytics,” data scientists use past patterns to guess what is likely to happen, or how an individual will act, in the future.

All of us have been interacting with this type of artificial intelligence for years. Machine learning has been used to create GPS systems, to make translation and voice recognition much more precise, to produce visual digital tools that have facial recognition or filters that create crazy effects on Snapchat or Instagram. Amazon uses artificial intelligence to recommend books, Spotify uses machine learning to recommend songs, and schools use the same techniques to shape students' academic trajectories.

Fortunately—or not, depending on one's point of view—we're not at the point where humanoid robot teachers stand at the front of class. The use of artificial intelligence in education today is not embodied, as the roboticists call it. It may have physical components, like internet of things (IoT) visual or audio sensors that can collect sensory data. Primarily, however, educational

artificial intelligence is housed in two-dimensional software-processing systems. This is perhaps a little less exciting, but it is infinitely more manageable than the issues that arise with 3-D robots.

In January 2019, the *Wall Street Journal* published an article with a very provocative title: “Colleges Mine Data on Their Applicants.”¹¹ The article discussed how some colleges and universities are using machine learning to infer prospective students' level of interest in attending their institution. Complex analytic systems calculate individuals' “demonstrated interest” by tracking their interactions with institutional websites, social media posts, and emails. For example, the schools monitor how quickly recipients open emails and whether they click on included links. Seton Hall University utilizes only about 80 variables. A large software company, in contrast, offers schools dashboards that “summarize thousands of data points on each student.” Colleges and universities use these “enrollment analytics” in determining which students to reach out to, what aspects of campus life they should emphasize, and assessing admissions applications.

AI Applications

Figure 1 shows a summary of the different kinds of applications that currently exist for artificial intelligence in higher education. First, as I've discussed above, is institutional

use. Schools, particularly in higher education, increasingly rely on algorithms for marketing to prospective students, estimating class size, planning curricula, and allocating resources such as financial aid and facilities.

This leads to another AI application, student support, which is a growing use in higher education institutions. Schools utilize machine learning in student guidance. Some applications help students automatically schedule their course load. Others recommend courses, majors, and career paths—as is traditionally done by guidance counselors or career services offices. These tools make recommendations based on how students with similar data profiles performed in the past. For example, for students who are struggling with chemistry, the tools may steer them away from a pre-med major, or they may suggest data visualization to a visual artist.

Another area for AI use in student support is just-in-time financial aid. Higher education institutions can use data about students to give them microloans or advances at the last minute if they need the money to, for example, get to the end of the semester and not drop out. Finally, one of the most prominent ways that predictive analytics is being used in student support is for early warning systems, analyzing a wide array of data—academic, nonacademic, operational—to identify students who are at risk of failing or dropping out or having mental health

Figure 1. AI Applications in Higher Education



The promise of AI applications lies partly in their efficiency and partly in their efficacy. **AI systems can capture a much wider array of data, at more granularity, than can humans. And these systems can do so in real time.**

issues. This particular use shows some of the real advantages of artificial intelligence—big data can give educators more holistic insight into students’ status. Traditionally, an institution might use a couple of blunt factors—for example, GPA or attendance—to assess whether a student is at risk. AI software systems can use much more granular patterns of information and student behavior for real-time, up-to-the-minute assessment of student risk. Some even incorporate information such as when a student stops going to the cafeteria for lunch. They can include data on whether students visit the library or a gym and when they use school services. Yet while these systems may help streamline success, they also raise important concerns about student privacy and autonomy, as I discuss below.


Lastly, colleges and universities can apply artificial intelligence in instruction. This involves creating systems that respond to individual users’ pace and progress. Educational software assesses students’ progress and recommends, or automatically delivers, specific parts of a course for students to review or additional resources to consult. There are often called “personalized learning” platforms. I put this phrase in quotation marks because it has been sucked into the hype machine, with minimal consensus about what personalized learning actually means. Here I’m using the phrase to talk about the different ways that instructional platforms, typically those used in a flipped or online or blended environment, can automatically help users tailor different pathways or provide them with feedback according to the particular error they make. Learning science researchers can put this information to


long-term use by observing what pedagogical approaches, curricula, or interventions work best for which types of students.

Promise and Perils

The promise of AI applications lies partly in their efficiency and partly in their efficacy. AI systems can capture a much wider array of data, at more granularity, than can humans. And these systems can do so in real time. They can also analyze many, many students—whether those students are in a classroom or in a student body or in a pool of applicants. In addition, AI systems offer excellent observations and inferences very quickly and at minimal cost. These efficiencies will lead, we hope, to increased efficacy—to more effective teaching, learning, institutional decisions, and guidance. So this is one promise of AI: that it will show us things we can’t assess or even envision given the limitations of human cognition and the difficulty of dealing with many different variables and a wide array of students.

Given these possible benefits, the use of artificial intelligence is also being framed as a potential boom to equality. With the improved efficacy of systems that may or may not require as much assistance from humans or necessitate that students be in the same geographical location, more students will gain access to better-quality educational opportunities and will perhaps be able to network with peers in a way that will close some of the achievement gaps that continue to exist in education. Lastly is the promise of a more macrolevel use of artificial intelligence in higher education to make gains in pedagogy, to see what is most effective for a particular student and for learning in general.





The use of artificial intelligence in higher education also involves perils, of course.² One is the peril of adverse outcomes. Despite the intention of the people who develop and use these systems, there will be unintended consequences that are negative or that can even backfire. To avoid these adverse outcomes, we should take into account several different factors. One of the first to consider is the data that these tools draw upon. That data can vary in quality. It may be old and outdated. Or it may be focused on and drawn from a subset of the population that may not align with the students being targeted. For example, AI learning systems that have been trained on students in a particular kind of college or university in California may not have the same outcomes or reflect the same accuracy

for students in another part of the country. Or an AI system that was based on Generation X students may not have the same efficacy for native digital learners.

Another data aspect concerns comprehensiveness. Does the data include information about a variety of students? There has been much discussion about this recently in terms of facial recognition. Scholars looking at the use of facial recognition by companies such as Google, IBM, Microsoft, and Face++ have shown that in many cases, these tools have been developed using proprietary data or internal data based on employees. The tools are much more accurate for light-skinned men than light-skinned women or darker-skinned men. In one study, the facial recognition tools

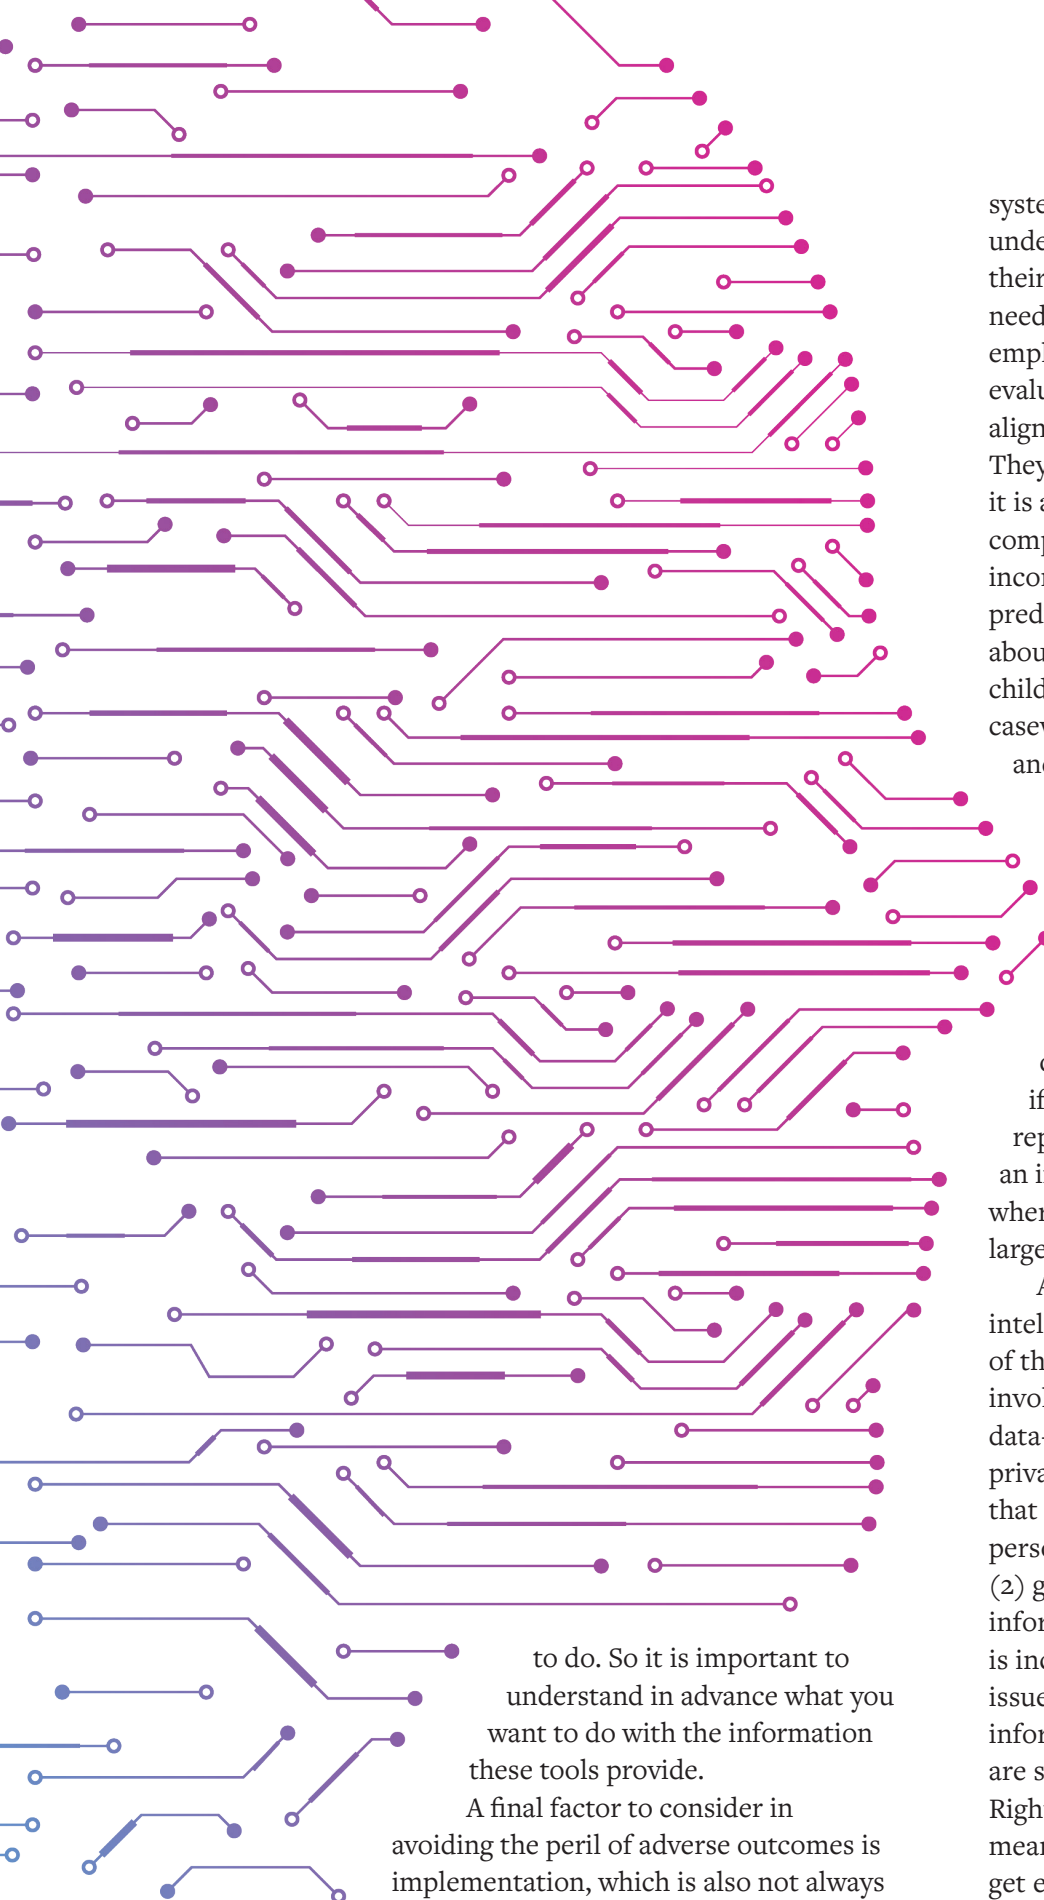
had nearly 100 percent accuracy for light-skinned men but only 65 percent accuracy for dark-skinned women. Joy Buolamwini, a co-researcher of this study, created her own, much more accurate tool simply by drawing from a broader array of complexion in the training data she used.³

Next to consider are the models that are created using this data. Again we face the issue of accuracy. Models are based on correlation; they are not reflective of causation. And as the [Spurious Correlations website](#) hilariously demonstrates, there are some wild correlations out there. Some correlations do seem to make intuitive sense, for example that people who buy furniture protectors are better credit risks, perhaps because they are more cautious. But the point of AI tools and models

is to show less intuitive, more attenuated correlations and patterns. Separating which correlations and patterns are accurate and which are simply noise can be quite difficult.

Algorithmic bias plays a role here. This is a real concern because it is something that can occur in the absence of discriminatory intent and even despite efforts to not have different impacts for different groups. Excluding a problematic or protected class of information from algorithms is not a good solution because there are so many proxies for things like race and gender in our society that it is almost impossible to remove patterns that will break down along these lines. For example, zip code often indicates race or ethnicity. Also, because artificial intelligence draws from existing patterns, it reflects the unequal access of some of today's current systems. A recent example is Amazon's hiring algorithm, which was criticized for being sexist.⁴ There is no evidence that Amazon had any intention of being discriminatory. Quite the contrary: Amazon used artificial intelligence to detect those characteristics that were most indicative of a successful employee, incorporated those characteristics into its algorithm, and then applied the algorithm to applicants. However, many of Amazon's successful employees, currently and in the past, were men. So even without any explicit programming, simply the fact that more men had been successful created a model skewed toward replicating those results.⁵

An additional, often overlooked factor in adverse outcomes is output. Developers' decisions shape how the insights that AI systems offer are instructed and interpreted. Some provide detailed information on various elements of students' learning or behavior that instructors and administrators can act on. Other observations are not as useful in informing interventions. For example, one predictive analytics tool estimated that 80 percent of the students in an organic chemistry class would not complete the semester.⁶ This was not news to the professors, who still wondered what



systems responsibly, teachers and staff must understand not only their benefits but also their limitations. At the same time, schools need to create very clear protocols for what employees should do when algorithmic evaluations or recommendations do not align with their professional judgment. They must have clear criteria about when it is appropriate to follow or override computer insights to prevent unfair inconsistencies. Consider the use of predictive analytics to support decisions about when caseworkers should investigate child welfare complaints. On the one hand, caseworkers may understand the complex and highly contextualized facts better than the machine. On the other, they may override the system in ways that may reflect implicit bias or have disparate outcomes. The people using these systems must know enough to trust—or question—the algorithmic output. Otherwise, they will simply dismiss the tools out of hand, especially if they are worried that machines may replace them. Good outcomes depend on an inclusive and holistic conversation about where artificial intelligence fits into the larger institutional mission.

A second peril in the use of artificial intelligence in higher education consists of the various legal considerations, mostly involving different bodies of privacy and data-protection law. Federal student-privacy legislation is focused on ensuring that institutions (1) get consent to disclose personally identifiable information and (2) give students the ability to access their information and challenge what they think is incorrect.⁷ The first is not much of an issue if institutions are not sharing the information with outside parties or if they are sharing through the Family Educational Rights and Privacy Act (FERPA), which means an institution does not have to get explicit consent from students. The second requirement—providing students with access to the information that is

to do. So it is important to understand in advance what you want to do with the information these tools provide.

A final factor to consider in avoiding the peril of adverse outcomes is implementation, which is also not always covered in the AI debates in the news or among computer scientists. To use these

being used about them—is going to be an increasingly interesting issue.⁸ I believe that as the decisions being made by artificial intelligence become much more significant and as students become more aware of what is happening, colleges and universities will be pressured to show students this information. People are starting to want to know how algorithmic and AI decisions are impacting their lives.

My short advice about legal considerations? Talk to your lawyers. The circumstances vary considerably from institution to institution.

Ethical Questions

Ethical questions revolve around consequences in terms of different groups and subgroups, educational values, and how AI systems might alter those values.

The Black Box

Unpacking what is occurring within AI systems is very difficult because they are dealing with so many variables at such a complex level. The whole point is to have computers do things that are not possible for human cognition. So trying to break that down ends up creating very crude explanations of what is happening and why.

Invisible Infrastructure

By choosing the variables to be fed into admission systems or financial aid systems or student information systems, these AI tools are creating rules about what matters in higher education. This leads to an invisible infrastructure. None of this is explicitly considered by the people implementing the infrastructure. The best example is when learning software specifies particular learning outcomes. That is, in essence, a high-core aspect of educational and institutional policy. But educators often overlook that fact when they adopt technology, not understanding that doing so is in some ways the equivalent of imposing an entirely different rubric, instead of standards, in the academic attainment.⁹

Authority Shifts

The entity doing the data collection and visualization is often a private company. That company is thus in charge of many decisions that will have an important impact and that will alter core values of systems in a way that is, again, not always visible. These private companies may be less directly accountable to stakeholders of the educational institutions—in particular, stakeholders such as students. It is important for us to consider this authority shift, and the shift in incentives, when using these technologies.

Narrowly Defined Goals

Applications that are based on data often promote narrowly defined goals. That is because in order to work, these systems must literally codify the results that are deemed optimal. This leaves less flexibility than is currently the case with human interactions in classrooms and on campuses. An example is acquiring an education broadly versus learning more narrowly. Optimizing learning outcomes—for example, additional skills acquisitions or better grades or increased retention—may crowd out more abstract educational goals promoting citizens capable of self-governance or nurturing creativity. The latter are aspects that one could technically, perhaps, represent in data, but doing so involves crude proxies at best. As a result, they may not be measured or prioritized.

Data-Dependent Assessment

Data-dependent assessment raises similar issues. Tools that collect information, particularly based on online interactions, don't always grasp the nuances that teachers might see in person. Consider the case where a student answers a question incorrectly. A machine will record a wrong answer. An instructor, however, may discount the error if she notices, for example, that the student clearly has a bad cold.

Divergent Interests

A divergent interest is sometimes between technology developers and institutions and sometimes between institutions and students. In the first instance, technology developers have an incentive to develop systems that use more and more data to get results that the developers can claim are more and more accurate. This allows them to show that their systems are making a difference. That may sometimes result in a rush to market or an emphasis on scale—which may not mean that the best-quality platforms are being used or that their efficacy is being assessed in any meaningful terms. This is certainly not true for all technology developers, but it is important to note.

More significant, and less obvious, is the divergent interest between institutions and students. The use of predictive analytics and early warning systems is often touted as a way to promote student retention by drawing attention to struggling or at-risk students. That is fine if the college or university is then going to institute intervention to try to ameliorate or prevent that outcome. But doing so is not always in the institution's administrative interest. In a famous example from a couple years ago, the president of Mount St. Mary's University, in Maryland, administered a predictive analytics test to see which students were most at risk of failing. The idea was to encourage them to drop out before the university was required to report its enrollment numbers to the federal government, thereby creating better retention numbers and improving its rankings. According to the president, his plan promoted the institutional interests for better statistics and was also in the students' best interest by preventing them from wasting money on tuition.¹⁰ Clearly, this goes into deeper questions of what the institutional and educational enterprise is and should be.

Elements to Consider and Questions to Ask

Several elements need to be considered to

ensure that the implementation of AI tools is optimal and equitable:

- **Procurement.** Pay close attention to the technologies and companies that will be most applicable to your particular student body in terms of the contractual obligations to provide data about your students. Make sure that if problems arise, you have contracted with a company that will be responsive to your problems.
- **Training.** Prepare those people who are going to implement and use these tools, and train them in the benefits and shortcomings of the tools.
- **Oversight.** Put in place a continuous process of examining whether the tools are working, whether they are more effective for particular groups of students, and whether they may be giving better numbers but not better outcomes. This is something that is difficult to do but is very important, because these tools can get outdated quickly.
- **Policies and Principles.** Create institutional policies surrounding the implementation of tools that rely on analytics, and cultivate principles that translate those policies into operational steps and actions.
- **Participation.** Get students' and faculty members' input about their concerns and what they would like to see from these systems. This step is often overlooked because it is messy and can lead to some controversy, but it generally creates a better result in the long run.

Regarding policies and principles, some of the best I have seen were developed in 2015 by the University of California Educational Technology Leadership Committee.¹¹ The committee listed six principles, elaborating on each: ownership; ethical use; transparency; freedom of expression; protection; and access/control. In addition, the committee recommended learning data privacy practices that security

providers can implement in the areas of ownership, usage right, opt-in, interoperable data, data without fees, transparency, service provider security, and campus security.¹²

Finally, to be successful, anyone considering an AI implementation within higher education should ask six essential questions:

1. What functions does the data perform? You can't just see a red, green, and yellow light about student success and take that at face value, at least not if you are the one implementing the systems and you want to do so responsibly.
2. What decisions don't we see? These are decisions not just about the computer processing but also about the categorization and the visualization.
3. Who controls the content? Is it you, or is it the technology provider? How comfortable are you with that? How comfortable are your professors with that?
4. How do we check outcomes in terms of efficacy, in terms of distribution, and in terms of positive and negative outcomes?
5. What gets lost with *datafication*? I use this word to describe doing these things based on data as opposed to on interpersonal or bureaucratic systems.
6. What—and whose—interests do we prioritize?

There are no easy answers, but asking these questions will give you a template for considering the less obvious aspects of these systems.

Conclusion

My final message? Do not surrender to the robot overlords just yet. Keep in mind that for all the hype and buzz, these AI tools are just computer systems. They can go wrong. They are created by humans. Their values are shaped by companies and institutions. Their data is not neutral but is defined by the historical patterns. Be cautious and

thoughtful about what you are doing with artificial intelligence, and remember: it's not magic. ■

Notes

1. Douglas Belkin, "[Colleges Mine Data on Their Applicants](#)," *Wall Street Journal*, January 26, 2019.
2. Manuela Ekowo and Iris Palmer, "[The Promise and Peril of Predictive Analytics in Higher Education: A Landscape Analysis](#)," New America policy paper, October 24, 2016.
3. Joy Buolamwini and Timnit Gebru, "[Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification](#)," *Proceedings of Machine Learning Research*, vol. 81 (2018).
4. Rachel Goodman, "[Why Amazon's Automated Hiring Tool Discriminated against Women](#)," *American Civil Liberties Union (ACLU) blog*, October 12, 2018.
5. Jeffrey Dastin, "[Amazon Scraps Secret AI Recruiting Tool That Showed Bias against Women](#)," Reuters (website), October 9, 2018.
6. Dian Schaffhauser, "[The Rocky Road of Using Data to Drive Student Success](#)," *Campus Technology*, July 26, 2018.
7. Elana Zeide, "Student Privacy Principles for the Age of Big Data: Moving Beyond FERPA and FIPPs," *Drexel Law Review*, vol. 8 (2016).
8. David Sallay and Amelia Vance, "[11 Steps to Privacy: The Right to Inspect and Review under FERPA](#)," *FERPA/Sherpa* (blog), April 24, 2019.
9. Elana Zeide, "The Structural Consequences of Big Data-Driven Education," *Big Data*, vol. 5 (2017).
10. Scott Jaschik, "[Are At-Risk Students Bunnies to Be Drowned?](#)" *Inside Higher Ed*, January 20, 2016.
11. Jim Phillips and Jim Williamson, "[UC's Learning Data Privacy Principles Gaining National Attention](#)," *UC IT Blog*, University of California, January 30, 2019.
12. See "[University of California: Learning Data Privacy Principles and Practices](#)" (website), accessed June 17, 2019.

Elana Zeide is Assistant Professor at the University of Nebraska College of Law. At the time of publication, she was a PULSE Fellow in Artificial Intelligence, Law, and Policy at UCLA School of Law for 2018–2020.

© 2019 Elana Zeide. The text of this article is licensed under the [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](#).

This article was originally published in *EDUCAUSE Review* on [August 26, 2019](#).

WAYS AI CAN HELP STUDENTS WITH DISABILITIES

This article was excerpted and adapted from an *EDUCAUSE Exchange* podcast (Gerry Bayne, producer): Judy Brewer, Carly Gerard, and Mark Hakkinen, [“The Impact of AI on Accessibility,”](#) *EDUCAUSE Review*, November 4, 2020.

In 2015 and 2016, nearly 20 percent of undergraduate students in the United States reported having a disability.¹ The real percentage is likely higher, given that many students choose not to disclose disabilities to their institutions. Their dropout rates are substantially higher and their graduation rates are significantly lower than these rates for nondisabled students.

Students with disabilities experience educational barriers that many other students do not, and they can have both visible and invisible needs. Artificial intelligence (AI) is being explored to improve and create tools for more accessible learning environments. Here are three ways AI can help these students.

1. Accessibility in Testing

Advanced speech synthesis technologies, which are based on machine learning models, are among the more promising applications of AI for students who rely on assistive technologies.² The quality of synthetic speech is becoming more natural and improving rapidly. For example, Educational Testing Service (ETS) used technologies from Amazon to replace some human recorded audio with synthesized speech for some supplemental test content. ETS improved the user experience for students with disabilities by reducing the turnaround time for producing alternate format materials and providing a more natural and clear text-to-speech voice for these students.

2. Content Descriptions

Turnaround time is significant when producing things like text descriptions or a complex set of test questions for students who are legally blind or have low vision. AI techniques could be used to automatically describe images. AI-based systems could also be used to do a “first pass” at describing content. Subject matter experts could then refine the content or, depending on the quality of the description, determine whether the content should be written from scratch.

3. Webpage Interactions

AI-based tools can also be used to help with interactions by people who are unable to see content. Tools like Apple Siri and Amazon Echo and Alexa provide ways of interacting with content through a spoken dialogue model. But there are many ways for AI features to expand. A “seeing” AI, for example, could help students who find the contents of a webpage to be too visually stimulating. Students could ask the virtual assistant to read aloud the headings on a page, allowing them to get a sense for how the page is structured, figure out where to go on the page, or skip content that is not relevant. Building this type of accessibility into the system that everyone uses—so that it simply comes onboard with every smart device—may also reduce the stigma

(and possibly the cost) associated with having to purchase separate accessibility tools or apps.

Challenges and Possibilities

AI-based design and development is often driven by the needs and behaviors of the “average user,” and from a user experience design perspective, people with disabilities typically fall outside of the usual experience. Automatic speech recognition (ASR) systems, for example, typically are optimized around common speech patterns, not around the speech patterns of people with speech disabilities. As a result, students who rely on ASR systems are more likely to be disadvantaged in educational and work settings where the ASR may not be optimized for them.

On the other hand, AI also holds great promise for people with disabilities. In the future, ASR systems may provide error-free closed-captioning rather than approximations. AI may also allow people with disabilities to fully control their environments—not only at home but also in the classroom and the workplace.

Full-scale automation may not yet be practical, but progress is being made. Some organizations are already using AI to assess conformance to accessibility guidelines. As this use becomes more widespread, conformance assessment will become more scalable. And as this use continues, we will find many other ways in which AI can be used to improve accessibility and ensure that students with disabilities have access to rich learning opportunities. ■

Notes

1. National Center for Education Statistics, Institute of Education Sciences, [“Postsecondary Education,”](#) chapter 3 in *Digest of Education Statistics: 2019* (Washington, DC: NCES, IES, U.S. Department of Education, 2021).
2. For more information about the technology needs of students with disabilities, see Dana C. Gierdowski and Joseph D. Galanek, [“ECAR Study of the Technology Needs of Students with Disabilities, 2020,”](#) *EDUCAUSE Review*, June 1, 2020.

© 2022 EDUCAUSE. The text of this work is licensed under a [Creative Commons BY-NC-ND 4.0 International License](#).

Combining AI with Human Oversight: 3 Benefits for Online Proctoring

By Jordon Adair

Certain aspects of how online education and assessments are delivered can greatly benefit from AI and machine learning. One of those aspects is online proctoring. Using “smart” technologies, institutional leaders in higher education can make human test monitoring more efficient, improve teaching and learning, provide better exam insights, streamline review processes, and build students’ confidence.


Today, online proctoring is accomplished in multiple ways: AI-only monitoring (fully automated), human-only monitoring, or a combination of both. Because each exam is different and there are many complex and unforeseeable situations that may occur, the most effective online proctoring model combines AI, machine learning, and human intervention to detect and prevent cheating and improve the entire testing experience for students *and* instructors.

Here are three ways that AI combined with human oversight benefits online proctoring.

1. Creates a Less-Invasive Experience for Students

Fewer distractions during exams means that students can focus on showing what they know. But some proctoring AI is overly sensitive, which can trigger flags and create unnecessary interruptions. For example, students who talk to themselves while working through a question can trigger a flag.

Most proctoring AI can detect sound, but some platforms have AI with smart *voice detection*. While these terms are often used interchangeably, how they impact the test



DEN RISE /SHUTTERSTOCK

experience is very different. *Sound detection* can trigger a flag for basically any sound, such as a dog barking or a cough. *Smart voice detection* triggers for specific keywords or phrases, such as “Hey Siri” or “OK Google.” Instructors can add custom lists for the AI to listen for, such as “What’s the answer?”

Having AI in place can quickly monitor and flag actions, which saves time, but some situations can make things tricky. For example, “OK Google” may be a keyword that triggers a flag. But instead of searching Google, a student could be saying, “Okay, Google was founded in 1998, so Yahoo is older.”

In this situation, combining a human proctor with the AI can prevent unnecessary interruptions. When human proctors get an alert from the AI, they can review the analysis window and see that the student wasn’t attempting to cheat, so there’s no need to intervene and disrupt the student.

2. Makes AI Smarter and Streamlines Instructors’ Reviews

Data such as potential violations and suspicious behavior, feedback from human proctors, and input from instructors can be collected by the proctoring platform and fed back to the AI. All

of this data helps the AI learn and adapt so as to provide educational institutions and instructors with better insights into the testing experience.

Since this continual feedback loop may be different for each institution, humans can customize the proctoring platform based on factors that are both important and unimportant to instructors. This feedback can also allow the AI to make better correlations between behaviors that may indicate true academic dishonesty and those that are false flags.

As the AI learns, it can adjust to help reduce unimportant or false flags due to oversensitivity. This helps both proctors and instructors focus on important situations related to academic dishonesty, and it streamlines instructors’ reviews. The combined power of both the AI and the human aspect creates a more equitable environment for students, who can feel confident that they won’t be inaccurately flagged.

3. Enables Nontraditional Exam Activities and Formats

AI is great for most traditional testing situations, but what if instructors want to test through a demonstration instead of a traditional format, such as multiple-choice questions? Using AI-only can overstep and create a frustrating experience for students and instructors. In these cases, combining AI with human input is crucial for protecting academic integrity and the students’ experience during complex assessments.

AI monitors for common anomalies such as two people in the room, unexpected voices or movements, restricted browsers, and leaked test questions. But in a nontraditional format or complex assessment, some of these activities may be allowed, and the AI needs to be told how to act.

SPONSORED CONTENT



Before a nontraditional exam starts, instructors can adjust test settings and decide which proctoring features to use. They can also give the human proctor instructions to further customize what is allowed during the exam and to provide accommodations for specific students.

For example, a nontraditional exam may ask students to complete a math problem using pen and paper. The instructor would ask the human proctor to turn on “Scratch Paper Allowed” so that students won’t be flagged for looking down at papers. Instructors can require students to use a sheet of blank white paper and a black or blue pen and to show the paper before and after they complete the problem. This ensures that students can complete the exam without interruption and cuts down on the need for human proctors to review flags.

While AI is a powerful tool that can improve online proctoring and the overall testing experience, it isn’t always perfect and shouldn’t be used alone. Instead of being a singular way to protect academic integrity, AI should be harnessed to complement, not replace, the efforts of human proctors and instructors. The combination of AI and human oversight is the key to creating an ideal testing experience for both instructors and students. ■

Jordan Adair is Vice President of Product at Honorlock.

© 2022 Honorlock.

This article was originally published in *EDUCAUSE Review* on [June 1, 2022](#).

Preventing
a Winter of
Disillusionment:

Artificial Intelligence and Human Intelligence in Student Success

By Linda Baer, Amanda Hagman, and David Kil



Student success,

in its various forms, is a top issue in higher education. Over the last decade, colleges and universities have worked to consolidate mountains of data into insights that can empower academic professionals to influence student success. Yet this cannot be accomplished using only human intelligence (HI). To facilitate an impact on student success, many institutions have employed artificial intelligence (AI) to help process and analyze data. AI, embedded in data systems, can allow institutions to better gather high-value data, monitor and uncover predictive risk indicators, and proactively respond to student behavior to promote student success.

Despite the high capabilities of these systems, they cannot be sustained outside professional HI, which gives meaning and direction to data insights. By providing enhanced information, AI helps humans to focus on insights relevant for student success impact and to proactively support student success. The promises of AI—that is, predictive models that create early alerts or evaluative tools to estimate the impact of interventions on student success—are possible only when HI and AI work together.

In “Student Success: 3 Big Questions,” Kathe Pelletier focused on what student success means, how it is measured, and whether or not student success is a mission-critical component of higher education institutions.¹ These are important foundational questions for improving student success. Next steps must address how leaders can build smarter student success models that scale and achieve sustainable results. This cannot be done without increasing the synergy between AI and HI.

Linking smart machines with human insight creates student success models that maximize outcomes while minimizing risk. As Diana Oblinger explains: “Machine learning allows computers to ‘consume’ information such as medical records, financial data, purchases, and social media and then develop predictions or recommendations. . . . These machines can create their own guidelines and discover patterns invisible to humans.” She quotes Garry Kasparov, the former world chess champion, who observed: “Humans are not being replaced by AI, we are being promoted. Machine-generated insights add to ours, extending our intelligence in the way a

telescope extends our vision. Think of AI as ‘augmented intelligence.’ Our increasingly intelligent machines are making us smarter.”²

Research on what contributes to student success and the growing focus on data and analytics set the stage for improving the ability to increase student retention and completion. We know more about student behavior and the activities that lead to success or risk. AI brings results to decision makers in real time. Predictive models allow discernment about which factors contribute to individual students’ progress and momentum. By combining student segments with learning life cycles, higher education professionals can align learner, time, and interventions into a model to maximize student success and decrease risk.

New technologies support the data mining, reporting, evaluation, and action by decision makers. As Heath Yates and Craig Chamberlain have noted, machine learning allows the modeling and extracting of useful information from data: “Adopting a machine learning–centric data-science approach as a tool for administrators and faculty could be a game changer for higher education.”³ Creating space for a synergistic relationship between HI and AI will be transformative.

But we face an obstacle: a winter of disillusionment. This can happen when AI hype leads to disappointment and criticism due to little-to-no tangible benefits. In fact, two AI winters have already occurred, in the 1970s and again in the 1980s.⁴ How can we prevent another such winter related to student success? Doing so requires that we become successful at improving student success, measured in a scientifically rigorous manner, by maximizing the symbiosis between HI and AI.

Defining AI and HI for Higher Education Objectives

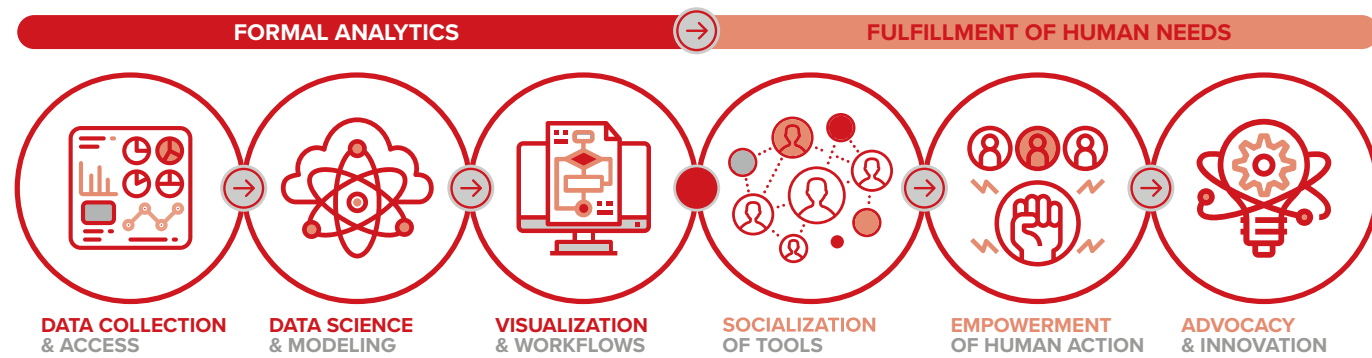
Data science is a discipline of constructing an intelligent system that ingests data from multiple sources, performs data transformations, and deploys various



machine learning algorithms in an attempt to make the system adapt and become more intelligent over time in solving business problems. Data science has greatly benefitted higher education by federating formerly siloed data, transforming the data into a useful state, and analyzing the data to identify insights that were previously hidden from view or took too much time to be of use for active students. Insights from data science efforts have included robust descriptions of student populations, predictive models, and even analyses to estimate the causal inference between institutional operations and key outcomes of student success.

AI refers to a system’s ability to interpret data correctly, learn from it, and achieve specific business goals through the judicious use of collected knowledge over time. Machine learning consists of a set of statistical and deep-learning algorithms that facilitate meaningful learning from data. AI uses automated logic and reasoning to streamline vast quantities of digital data and automatically improve knowledge over time.

Figure 1. The Lifecycle of Sustainable Analytics



Source: Mitchell Colver, “The Lifecycle of Sustainable Analytics: From Data Collection to Change Management,” unpublished paper, Office of Student Analytics, Utah State University (Logan, UT, 2018). Reprinted with permission.

Unfortunately, AI, due to its dependence on learning from data, cannot think outside the box, meaning that making open-set decisions based on new patterns in data can be very challenging without HI. For example, mortgage-backed security pricing algorithms blew up in 2008 because they were trained on the previous three years of data—a time when home prices had been rising.⁵ Furthermore, intentional intervention design can benefit from (1) human creativity in integrating knowledge from descriptive, predictive, prescriptive, and impact analytics, and (2) deep understanding of behavioral science, which is often missing in quantitative institutional data. That is, while AI is good at chewing through a large volume of data to find patterns and make predictions, piecing everything together for coordinated actions and student success outcomes still requires HI. This is the essence of the synergy between AI and HI.

Since the beginning of time, logic and reasoning have been the hallmarks of HI: people analyze and interpret the perceived variables within their environment. Unfortunately, the number of perceived variables has exploded with the accumulation of digital data. Colleges and universities are awash in data from students’ participation in almost every aspect of campus life. Higher education professionals have access to far more data than they can interpret and utilize to influence student success. Fortunately, AI can assist HI in

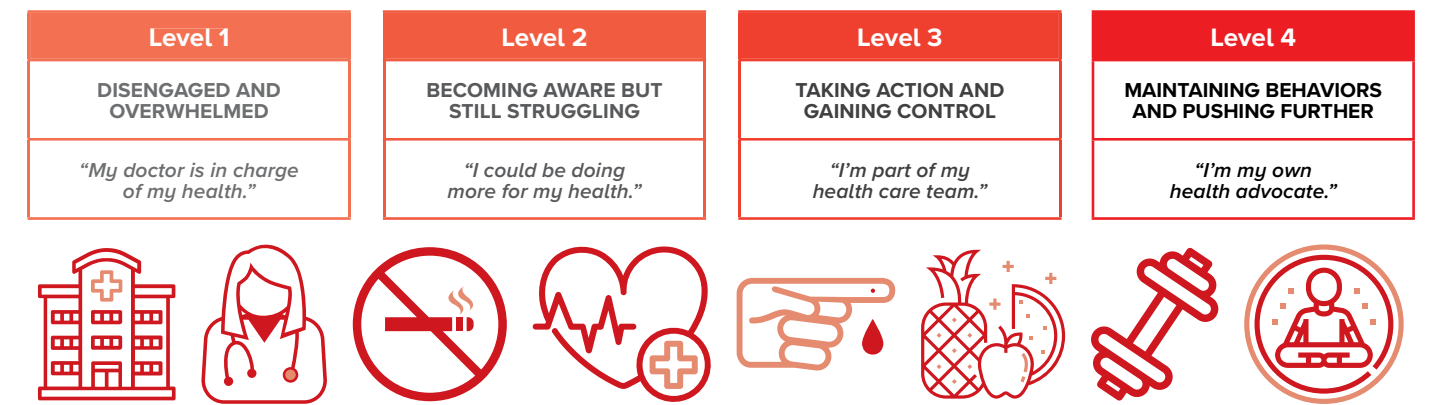
processing and organizing insights that historically have been hidden from view. Working together, AI and HI can leverage insights from data to directly influence student success and institutional functions.

A useful model for understanding the relationship between AI and HI is “The Lifecycle of Sustainable Analytics” (see figure 1).⁶ This integrated model acknowledges the necessity of AI and HI to solve 21st-century problems in higher education. The model makes a distinction between the steps in formal analytics (data collection, data science, and visualization) and the steps in the fulfillment of human needs through analytics (socialization, empowerment, and advocacy). Any data initiative must be socialized to cover not only the how of using AI insights but also the why and when of using these insights. Higher education professionals must understand how AI promotes them and complements their work so that they can feel empowered to incorporate AI technologies into their daily actions. Finally, professionals must see how the insights can be used to advocate and innovate in their work. Finding a harmony between AI and HI is necessary for the success and sustainability of data science initiatives.

Lessons from Health Care

As higher education adopts AI methods to assist HI in the immense task of student success, we can learn from fields that pioneered AI methods to tackle complex

Figure 2. A Portfolio-Driven Approach to Patient-Care Optimization



Source: Linda Baer, Amanda Hagman, and Dave Kil, “What Leaders Need to Know about Scaling Student Success Programs,” WCET 31st Annual Meeting, Denver, CO, November 5, 2019. Reprinted with permission.

problems. An early leader of AI in industry was the health-care system. For example, in 2004 one health-care company built a patient-risk predictive model that outperformed the industry-standard model by over 20 percent. The company then developed a lifestyle coaching program that incorporated salient behavioral science and patient-activation principles. The company ran a pilot program on the diabetic population, measured outcomes, and found statistically significant positive results. Everyone was happy, and the company decided to expand the program to all patients.⁷

When the company measured outcomes again, however, they were very surprised to find negative outcomes: feedback from health coaches indicated that the patients who received outreach were much sicker than the initial pilot population. Instead of giving up, the company decided to dig deeper. Drill-down impact analysis showed that although some patient segments, such as those with diabetes or cardiovascular diseases, benefited from lifestyle coaching, patients with far more serious conditions and comorbidities did worse. Analysis of patient-coach interaction data, along with coach-level impact analysis, soon revealed that there was no one-size-fits-all intervention program.

These findings, along with strong encouragement from the company’s executive team, led to a new, portfolio-driven approach to patient care, with programs

catering to specific needs of various patient segments (see figure 2). Furthermore, the company measured the impact of all patient-care programs monthly, reviewing the results and discussing opportunities for performance and process improvement in a monthly steering committee meeting attended by all senior executives, clinical-program owners, and data scientists. This is a clear example of HI-AI synergy that led to a systemwide improvement in outcomes.

Implications for higher education from the health-care example are fascinating. First of all, making predictions is less important than knowing *how* to create a portfolio of programs personalized to population segments with specific needs. Predictions can help academic professionals focus on the right students, but knowing how to help them is the key here. Thus an important lesson learned from health care is to transform the AI and HI relationship from risk prediction to impact prediction. Impact predictions analyze how institutional programming is influencing student success across multiple student segments. Quantifying the impact of student initiatives allows the higher education institution to build a portfolio of student services. Drilling down into evaluations of the programs reveals what works and for whom and in which operational settings. In this process, higher education professionals will become equipped to prescribe programming that

can promote student success with existing resources. Campuses use a number of interventions to influence student success, but it is very difficult to improve without rigorously measuring their efficacy for continuous learning and portfolio optimization (i.e., resource allocation optimization given that everyone operates under a finite amount of resources).

Learning from pioneering health-care companies, higher education must foster the working relationship between AI and HI. Although many higher education institutions have adopted AI analytic systems, a report jointly produced by AIR, EDUCAUSE, and NACUBO calls for a much stronger approach to the use of analytics in student success. It concludes: “With the change-making capacity of analytics, we should be moving aggressively forward to harness the power of these new tools for the success of our institutions and our students. However, so far higher education has failed to follow talk with decisive action.”⁸

Some colleges and universities have indeed reaped benefits in terms of student retention, but others have been underwhelmed with the productiveness of AI systems on their campuses. A major problem may stem from the belief that transformative changes should flow spontaneously from AI analytic insights, but this ignores the key role played by the HI of higher education professionals. One example is the low prioritization of professional development at some institutions that have adopted sophisticated AI systems.⁹ HI must be trained on how to take insights from AI systems and innovate practice to improve student success.

In short, the goal of AI in higher education is to help design and execute intentional interventions in order to maximize the probability of student success. This moves HI away from a focus on repetitive and uninspiring work and toward tasks that inspire and reward us. Of particular interest here is the Fogg behavior model, which talks about aligning core motivators,

simplicity factors, and behavior triggers to increase the likelihood of humans performing targeted behavior.¹⁰ AI simplifies what we need to know about students and existing programs so that we can put together an action plan with confidence of its utility. Such intentional intervention design work appeals to our core motivators, giving us pleasure in seeing the fruits of our creative and mission-driven work. Furthermore, understanding the right behavioral triggers for students to comply with carefully designed calls to action can lead to a virtuous cycle of higher compliance and better outcomes. That is, having an evidence-based intervention recommendation adds to simplicity and appeals to core motivators, leading to improved odds of designing intentional interventions and impact success.

The building blocks for this transition between prediction and impact must include AI and HI working together toward the following:

1. Understand who is at risk, why, and what can move the needle on student success
2. Organize existing data and evaluate the need for improved data-capturing
3. Audit current programming and initiatives using impact analyses to discover what is working and for whom
4. Match at-risk students with programs shown to influence student success for similar students
5. Create evidence-based student success knowledge with learning lifecycle management and continuous evaluation as programs are adjusted to reflect intervention insights
6. Develop an action plan from evidence-based intervention data and evaluate results¹¹

Leveraging the benefits of AI and HI initiatives requires the above building blocks. Jonathan Zittrain has explored the pernicious nature of intellectual debt associated with AI when we do not know

how something works; failing to consistently train HI to understand and leverage insights from AI systems creates this intellectual debt.¹² At Utah State University, the Center for Student Analytics has taken on the task of empowering professionals to utilize insights from AI as a way to innovate university practices for improving student success. This has been accomplished by fostering a positive relationship between HI and AI and by helping professionals to see how these modern tools promote their current practices. The Center for Student Analytics at Utah State University has also established professional training as an institutional priority. Instead of receiving mere point-and-click training, professionals discover how to leverage insights from analytics into daily practices. They also learn about professional intentionality and the ethics of using big data in higher education. Dedicating resources to the empowerment of university professionals with modern technology has proven a boon to the culture of innovation within the institution.

Combined HI and AI in Action

What does combining AI and HI mean for student success models? Currently, smarter student success is possible by balancing AI and HI. Thanks to improved insights from AI, HI can concentrate on which actions and interventions provide the most impact for students.

Grinnell College has leveraged this balance between AI and HI by addressing the science of intervention to provide faculty and staff with information on its effectiveness. In “Blending Human Intelligence and Analytics for Student Success,” Randall J. Stiles and Kaitlin Wilcox state: “Colleges and universities have long relied on human-intelligence networks made up of faculty, professional advisors, other administrators, and students themselves to find the best balance of challenge and support for individualized learning and to monitor student progress.” Staff at Grinnell have

integrated learning analytics with HI networks “so that alerts, predictive models, and outreach to students might be improved.”¹³ This blending was based on the work of Thomas H. Davenport and Julia Kirby, who talk about *augmentation*, defined as “starting with what minds and machines do individually today and figuring out how that work could be deepened rather than diminished by a collaboration between the two. The intent is never to have less work for those expensive, high-maintenance humans. It is always to allow them to do more valuable work.”¹⁴

This cultural shift toward a balance between HI and AI can be seen in an example at Utah State University. A program designed to promote new freshmen’s integration into campus life was in jeopardy of losing funding. In the program, students attending academic and co-curricular programming accumulated points toward earning a monetary reward and a reception with executive-level university professionals. An impact evaluation revealed significant gains in student persistence for students who participated. Specifically, students who participated in the program were 2.7 percent more likely to persist than similar students who did not participate. This gain in persistence was associated with retaining an additional 38 students each year. The program was especially helpful for students who were most at risk of leaving the university.¹⁵

Given these insights—that (1) the program was effective and (2) it was influential for students at risk of leaving the university—the orphaned program was adopted by the Student Affairs Office. Unfortunately, while in transition, the program lost a large portion of its funding. In response to the decreased funding, university professionals reflected on their experience with the program (HI) and investigated the data (AI). In a facilitated discussion with the data team, university professionals added their contextual insights (HI) to the data. One HI insight revealed that many students were very eager to receive the monetary reward. Staff thus decided to keep the

Without HI, the AI technologies will fall short of our expectation of improved student success. Colleges and universities need to expand their capacities in data technologies in tandem with expanding their human capacities to ingest, incorporate, and innovate.

monetary reward for participation but cut the reception with university executives.

The following semester, the program was evaluated again with an impact analysis. Interestingly, the removal of the reception resulted in a reduced impact, from the 2.7 percent increase in persistence to a 1.1 percent increase in persistence. In other words, this programmatic change shifted from retaining 38 students a year to only 14. While anecdotal evidence from the first round of evaluation suggested the monetary reward was the largest motivator, losing the reception hurt the program. Unfortunately, the programmatic budget was not changed. Instead, university professionals worked within their constraints to identify no- or low-cost alternatives to the reception. They were able to pull together enough resources for several raffle drawings for meal plans, parking passes, and other university goodies. The impact of this change is not yet known, but the program is on track for an evaluation this spring. Regardless, one thing is clear: the university has established a cadence to quickly evaluate the impact of its programmatic changes. This symbiosis of AI and HI opens countless avenues for accountability, innovation, and advocacy for university programming.

Utah State University has also undertaken the task of evaluating existing student initiatives across campus using impact analyses with a common outcome of persistence. The sweeping project has given rise to a better description of how services are influencing student persistence. It is also uncovering insights about which students are benefiting from which initiatives. Through this process, students can be prescriptively matched to the initiatives that support their individual needs and success.

The most current example of this effort is the *Student Analytics Look Book*, a student-facing document that highlights analytical insights derived from predictive modeling and impact analyses of student initiatives.¹⁶ Promoting these insights through a *Look Book* to students and university professionals democratizes insights for the betterment of the student experience.

Given the above examples of HI-AI synergy, the desired output of AI systems is the knowledge base on how to improve business outcomes. As an analogy, the core mission of many precision medicine companies and nonprofit health organizations is to build the evidence-based treatment efficacy knowledge base as a function of a patient's clinical condition, treatment history, and molecular profile.¹⁷ In *What Works Clearinghouse (WWC)*, only twelve interventions in postsecondary education meet WWC guidelines for being a proven high-impact practice (as of November 13, 2019). Furthermore, these interventions have so many moving parts that scaling and replicating them at other institutions is very difficult, as well as very expensive, to implement. In addition, most colleges and universities are not consistently evaluating their implementation of the twelve WWC high-impact practices with impact analyses on a regular basis.¹⁸ In short, there is a strong moral imperative that we build the evidence-based student success knowledge base systematically in a scalable, cost-effective manner by fusing the most salient attributes from AI and HI.

Conclusions and Future Directions

David Watson recently lamented that while AI has been conceptualized in anthropomorphic terms, its true abilities have been vastly

overstated, robbing us of our own autonomy.¹⁹ Instead, as we have argued above, a balanced investment in AI technologies and HI capital can take AI tools to the next level. Without HI, the AI technologies will fall short of our expectation of improved student success. Colleges and universities need to expand their capacities in data technologies in tandem with expanding their human capacities to ingest, incorporate, and innovate.

Higher education has the power to prevent another AI winter of disillusionment related to student success. To ensure that the use of AI leads to tangible student success outcomes, we must champion the symbiosis between human intelligence and artificial intelligence. ■

Notes

1. Kathe Pelletier, "[Student Success: 3 Big Questions](#)," *EDUCAUSE Review* 54, no. 4 (Fall 2019).
2. Diana G. Oblinger, "[Smart Machines and Human Expertise: Challenges for Higher Education](#)," *EDUCAUSE Review* 53, no. 5 (September/October 2018).
3. Heath Yates and Craig Chamberlain, "[Machine Learning and Higher Education](#)," *EDUCAUSE Review*, December 18, 2017.
4. "[AI Winter](#)," Wikipedia, accessed December 10, 2019.
5. Felix Salmon, "[Recipe for Disaster: The Formula That Killed Wall Street](#)," *Wired*, February 23, 2009.
6. Mitchell Colver, "[The Lifecycle of Sustainable Analytics: From Data Collection to Change Management](#)," unpublished paper, Office of Student Analytics, Utah State University (Logan, UT, 2018).
7. Frances B. Shin and Dave Kil, "Integrated Wellness and Health Management: Leveraging Predictive Models and Resource Allocation Optimization," Third Annual Conference on Optimizing the Implementation of Predictive Modeling, Las Vegas, NV, June 2006.
8. Association for Institutional Research (AIR), EDUCAUSE, and National Association of College and University Business Officers (NACUBO), "[Analytics Can Save Higher Education. Really](#)" [2019].
9. Colver, "[The Lifecycle of Sustainable Analytics](#)"; Gerald C. Kane, Anh Nguyen Phillips, Jonathan R. Copulsky, and Garth R. Andrus, *The Technology Fallacy: How People Are the Real Key to Digital Transformation* (Cambridge: MIT Press, 2019).
10. B. J. Fogg, "A Behavior Model for Persuasive Design," Persuasive Technology, Fourth International Conference, Claremont, CA, April 26–29, 2009.
11. David Kil, Linda Baer, and Amanda Hagman, "Sherlock Holmes Redux: Putting the Pieces Together," chapter 7 in Colleen Carmean and Linda Baer, eds., *An Analytics Handbook: Moving from Evidence to Impact* (Ann Arbor, MI: Society for College and University Planning, 2019).
12. Jonathan Zittrain, "[The Hidden Costs of Automated Thinking](#)," *New Yorker*, July 23, 2019.
13. Randall J. Stiles and Kaitlin Wilcox, "[Blending Human Intelligence and Analytics for Student Success](#)," EDUCAUSE Learning Initiative (ELI) Research Brief, August 2016.
14. Thomas H. Davenport and Julia Kirby, *Only Humans Need Apply: Winners and Losers in the Age of Smart Machines* (New York: HarperCollins Publishers, 2016), p. 63.
15. Amanda Hagman, "Impact Analysis of the Passport Program on Student Persistence," unpublished paper, Utah State University, Logan, UT, 2019.
16. Center for Student Analytics, Student Insights Report, issue 1 (fall 2019), Utah State University, Logan, UT.
17. See, for example, "[Development and Update of the NCCN Guidelines](#)," National Comprehensive Cancer Network (website), accessed November 13, 2019.
18. J. Louviere, "Persistence Impacts on Student Subgroups That Participate in the High-Impact Practice of Service Learning" (PhD dissertation, Utah State University, Logan, UT, 2019).
19. David Watson, "The Rhetoric and Reality of Anthropomorphism in Artificial Intelligence," *Minds and Machines* 29, no. 3 (September 2019).

Linda Baer is Senior Consultant at Linda L. Baer Consultant and Senior Fellow at Civitas Learning.

Amanda Hagman is Director of Strategic Planning at dataRAE. At the time of publication, she was Data Scientist, Center for Student Analytics, at Utah State University.

David Kil is Chief Science Officer at prescienceAI. At the time of publication, he was Chief Scientist at Civitas Learning.

© 2020 Linda Baer, Amanda Hagman, and David Kil. The text of this article is licensed under the [Creative Commons Attribution 4.0 International License](#).

This article was originally published in *EDUCAUSE Review* on [May 18, 2020](#).

For Higher Ed, Chatbots Need to Be Smart

By Mary Frances Coryell

With the onset of the COVID-19 pandemic, chatbots quickly became a critical tool for a variety of operations, especially in colleges and universities. The strictly virtual interactions with students dramatically increased the need to be digitally nimble enough to produce optimal online experiences. Chatbots facilitate a very useful function as part of that experience, but only if they can adequately and reliably serve peoples' needs.

Many institutions that purchased templated chatbots in the past have come to regret that decision. Complaints about the bot's inability to sufficiently answer questions (or in some cases, causing more confusion) are commonplace. Anyone who has found themselves yelling "agent!" into their phones when speaking with a company's voicebot can likely understand the frustrations associated with an unhelpful chatbot.

Historical Issues

The primary challenge leading to this level of dissatisfaction is a chatbot's inability to truly understand what is being asked. The degree to which a bot can understand the intent of a user's question, as opposed to strictly relying on the "correct" combination of terms, is what determines whether the chatbot will be able to successfully resolve the user's issue.

Meanwhile, most stakeholders do not have the time, training, or inclination to develop an intelligent chatbot to address this problem. Their hope is that once deployed, the bot will have the ability to absorb the bulk of routine calls or emails to free up time for more meaningful aspects of their jobs. Unfortunately, this is seldom the case when the bot's responses are templated, requiring a great deal of time to manage the quality of its responses to real-world scenarios.

In higher education, IT help desks, along with admissions and financial aid offices, were among the earliest adopters of chatbot technology. Staff in these areas recognized the potential for chatbots to resolve a significant volume of questions before those questions became support tickets. However, they also learned that to work effectively, bots require ongoing development. Essentially, the problem comes down to resource constraints around building a sophisticated bot.¹

A Response to the Problem

A truly artificially intelligent chatbot that meets the standard of prestige required by higher education institutions needs to possess several characteristics.

Customizable

When an institution purchases a chatbot, the bot's knowledge must be customized to the institution and be distinctive from the knowledge of any other bot. The information it provides to users must be specific and complete with as few clicks as possible. The bot's knowledge should be curated straight from a variety of sources developed and managed by each institution.

Multi-Channel

Students are accustomed to multi-channel access, so chatbots must be accessible on a variety of channels. Students should be able to take the conversation beyond websites to SMS, social media, email, and voice assistants such as Amazon's Alexa and Apple's Siri.

Multilingual and Resistant to Implicit Bias

Chatbots must be inclusive and able to support diverse student populations. A chatbot that speaks more than one hundred languages can broaden the range of students it can help. In addition, a chatbot based on a range of user interactions can support a wide variety of users and eliminate implicit bias.

Secure and Scalable

Given the structure of higher education institutions, chatbots must be deployable in multiple departments, each with its own distinct knowledge and objectives. Administrators need a bot that does not require extensive development in order to provide high-quality, customized interactions in each area. A bot with technology that builds itself will put AI within reach across the institution and deliver anonymous, aggregate data that administrators can use to keep pushing the needle forward.

Self-Learning and Increasingly Intuitive

Finally, a chatbot must not be easily stumped or require an inordinate degree of effort to stay up-to-date. To that end, chatbots should not be exclusively rules-based or rely on keywords to understand meaning. Instead, a bot must learn from each conversation it has, perpetually improving its ability to answer questions at an increasing level of complexity and, ideally, to update its own knowledge when possible. As is the mission with any AI instrument, the chatbot should mimic a human resource and apply knowledge from its gained experience.

Today's Outcomes

Chatbots that incorporate all of these elements at once can have a significant impact on higher education. The bot implemented at Broward College, an institution with a student population of more than 65,000, reduced call volume by 9.6%—a total of 30,041 calls. Broward leaders estimate that this led to a cost savings of \$210,287. Their bot also deflected live chat conversations by handling 60,842 inquiries, leading to an additional cost savings of \$304,210. The University of Portsmouth reduced live chats by 50%, while Temple University reduced calls by the same margin. Meanwhile, with the help of "SoonerBot,"

the University of Oklahoma enrolled its largest class of first-year students in Fall 2019.²

These achievements were possible only because their chatbots kept learning, increasing the scope of support they could provide and delivering complete and accurate answers that fully satisfied the needs of students.

Next Steps

Those in higher education who have been seeking a chatbot solution to produce real, tangible results for their institution should take the following steps:

1. Get others at your institution onboard when beginning the process, to ensure that goals are aligned throughout deployment.
2. Evaluate the internal resources you want to commit to the ongoing success of the project (e.g., SMEs for each department, staff for live chat escalations).
3. Seek out a chatbot that fits your needs as well as those of your end users (e.g., breaking down silos, providing direct answers, rating chats).

As AI advances in the coming years, chatbots will get increasingly smarter, more intuitive, and as a result, more valuable for higher education. ■

Notes

1. For help selecting among the various types of chatbots on the market today, see [How to Select a Chatbot in 2022](#), Ivy.ai (website), accessed April 25, 2022.
2. See "[Case Studies](#)," Ivy.ai (website), accessed April 25, 2022; "[OU Uses Artificial Intelligence in Recruitment](#)," *Inside OU*, September 18, 2019.

Mary Frances Coryell is Chief Revenue Officer at Ivy.ai.

© 2022 Ivy.ai

This article was originally published in *EDUCAUSE Review* on [June 1, 2022](#).




Ivy.ai

NO CHEAT CODE

Exploring the Intersection of Empathy and Technology

By Tara Hughes



What one factor plays the most important role in making technology work in higher education?

For those of us who spend our days managing email servers and learning management system (LMS) integrations, “empathy” may not be the first thing that comes to mind. But the past two years have reinforced just how crucial empathy and emotional intelligence can be for technology leaders seeking to serve students at their institutions.

Let me explain. Before the COVID-19 pandemic, I ran the shared services department at California State University (CSU) Channel Islands. My role included building the knowledge base for Ekhhobot, our AI chatbot. After researching AI chatbots and learning that they can significantly reduce summer melt and boost retention, I thought my job would be simple: teach the bot what it needs to know, sit back, and let it rip.¹

I hoped that this technology would essentially act as a “cheat code” that could, without a lot of input from me or my team on specific content, help us with the difficult task of helping students navigate and complete their education. I was ready to play my part in supporting the ambitious [Graduation Initiative 2025](#), which aims to increase graduation rates and eliminate opportunity gaps for all students in the CSU system.

But a few months in, we weren’t getting the response we had expected. Instead of gratitude and student persistence, students were opting out of connecting with the chatbot—sometimes with a bit of blue language thrown in. (Sure, the name of our institution includes the word “islands,” but there’s no need to talk like a sailor!)

In January 2020, my team and I found ourselves in a pickle: the transactional relationship between our chatbot and our students wasn’t moving the needle on our engagement and retention goals. We didn’t realize that the nature of the communication itself was our roadblock.

We weren’t confident enough to be picky about what the chatbot said to students. As a result, Ekhhobot had no personality, and students were responding accordingly. But then the pandemic hit, and that’s when things changed.

By March 2020, we still hadn’t figured out the right way to communicate with students through the chatbot, but we knew we had to say something. We couldn’t leave our students hanging. So, we had the chatbot send all CSU Channel Islands students a silly meme. Suddenly, students began engaging in ways they hadn’t before. That led us to start experimenting with other ideas, such as knock-knock jokes and emoji smiley faces. We even built a Spotify playlist based on what students told the chatbot were their favorite songs at the time.

Underpinning all this work was a sense of empathy and compassion for the fact that our students were living through an incredibly challenging time. Sometimes that meant providing proactive information about the counseling center or other campus resources. Sometimes it meant sending an emoji or two. Regardless, the communication with students always came from a place that wasn’t transactional but rather was rooted in the genuine, face-to-face interactions that we know translate to a sense of belonging and motivate students to persist and work hard.

Students’ responses have provided the most powerful proof that an empathetic approach works. They treat Ekhhobot almost like a pet or a friend. They thank the bot for giving them advice, and they’re often comfortable acknowledging when they’re feeling stressed or anxious (which we can then elevate to campus counselors to provide one-on-one support). That’s pretty unique when you think about it. When was the last time you thanked the disembodied bot embedded in your phone?

Perhaps most importantly, our empathetic approach is helping us to make systemic changes across the institution. We’ve used Ekhhobot to survey students about what they like—and don’t like—about the remote learning experience. We then sent that feedback to our office of Teaching and Learning Innovations. More than one-third of students responded to the survey (a record number), and we had enough data to provide institutional leadership with a brief to help them double down on what was working and fix what wasn’t.

Growing up as an avid video game player, I always loved cheat codes, especially those that let you leapfrog whole sections of the game to get to the end. I had hoped that technology could do the same thing for our students and propel them to where we wanted them to be, with little or no input from any of us humans. The pandemic jolted me out of that belief, and it also helped me to understand something far more consequential: when technology is implemented thoughtfully and empathetically, the impact can be profound.

As we navigate our way toward a new normal, one thing is certain: Ekhhobot will keep sending students jokes, asking them for music recommendations, and helping them access the resources they need. The lesson we learned during the pandemic is that no one of those things is any more critical than another. Without the silly stuff, we can’t build the relationships that help students listen when it’s time to talk about more serious topics.

The experience of the past two years and our work to be more mindful about how we use edtech have been nothing short of transformative, and I hope other institutions can learn from our experience. If we think of technology as a tool to extend and amplify meaningful human interactions, students’ experiences will be all the better for it. ■

Note

1. Hunter Gehlbach and Lindsay C. Page, “Freezing ‘Summer Melt’ in Its Tracks: Increasing College Enrollment with AI,” *Brown Center Chalkboard* (blog), Brookings, September 11, 2018; Tim Renick and Lindsay Page, “What Does It Take for Nudging to Impact College Students’ Success?” *Higher Ed Dive*, September 4, 2020.

Tara Hughes is Interim Chief Information Officer at California State University Maritime Academy.

© 2022 Tara Hughes. The text of this work is licensed under a [Creative Commons BY-NC 4.0 International License](#).

This article was originally published in *EDUCAUSE Review* on [February 3, 2022](#).



EDUCAUSE

UNCOMMON THINKING
FOR THE COMMON GOOD

EDUCAUSE is a nonprofit association and the foremost community of IT leaders and professionals committed to advancing higher education.

EDUCAUSE Board of Directors

Helen Norris, Chair
Vice President & Chief Information Officer
Chapman University

Sharon Pitt, Vice Chair
Vice President for Information Technologies
and CIO
University of Delaware

Amelia Parnell, Secretary
Vice President for Research and Policy
NASPA - Student Affairs Administrators in
Higher Education

Celeste Schwartz, Treasurer
Vice President for Information Technology &
Institutional Effectiveness
Montgomery County Community College

Allan Chen
VP-IT/CTO
California Institute of the Arts

Elias Eldayrie
Vice President and Chief Information Officer
University of Florida

Gayleen Gray
Assistant Vice President and CTO
McMaster University

Tony Moore
Chief Information Officer
Prairie View A&M University

Carol Smith
Chief Information Officer
DePauw University

Cheryl Washington
Chief Information Security Officer
University of California, Davis

Melissa Woo
Executive Vice President for Administration
& CIO
Michigan State University

Sherri Yerk-Zwickl
Vice President for Information Technology
& CIO
Campbell University

Ex Officio
John O'Brien
President & CEO
EDUCAUSE

Strengthening Higher Education Together

EDUCAUSE thanks the sponsors of this Special Report on artificial intelligence.

