



Public University Boards and Artificial Intelligence

Report

Brian J.A. Boyd

U.S. Faith Liason
The Future of Life Institute

Executive Summary

Many technologies help us do one specific thing: a hammer and a boat, for example, have an obvious function. Other technologies are general-purpose and can be put to many uses: electricity, oil, and the internet are important because of their versatility. A few technologies, such as writing, are capable of directly making us better at making new technologies. Artificial intelligence is a unique general technology that can be pruned for specific-use cases, including the development of more powerful technology.

Universities, like other key industries and organizations, are being and will continue to be significantly affected by AI. Because institutions of higher education exist to curate and foster thought, they face special opportunities and risks from AI's potential to both amplify and undermine human thinking.

AI's rising prominence comes at a precarious moment for institutions of higher education. A demographic cliff of fewer college-age students threatens enrollment and therefore program health and campus budgets.¹ Widespread public skepticism about the value of college multiplies those effects.² Many colleges will likely close in the coming years, following hundreds that have already closed over the past two decades.³

The boards of public universities, the ultimate governing authority of these sprawling, essential institutions, must be prepared to navigate these challenging times. This report considers what boards of public universities can do to meet the moment.

First, we clarify the terms under consideration: What is AI, and what is the university?

About Us

The Manhattan Institute is a community of scholars, journalists, activists, and civic leaders committed to advancing economic opportunity, individual liberty, and the rule of law in America and its great cities.



Second, we take up a helpful image of the university offered by Niall Ferguson, “the cloister and the starship,” and expand it to illustrate the principles that should guide universities. If the humanities refocus on their core mission to form free and virtuous citizens, while the sciences and applied disciplines treat AI as a tool that could benefit the commonweal, then public universities can stay true to their long-standing mission while leading their states into the future.

Third, we look at specific policies for the broader university, including questions of academic integrity and AI-use disclosure.

Introduction: Setting the Terms

What Do We Mean by “AI”?

In January 2026, leading AI company Anthropic described the latest version of its chatbot Claude as “a genuinely novel kind of entity in the world” with “emotional states” and “existential questions” of its own.⁴ A more muted take on AI, which we will follow in this report, comes from Microsoft research scientist Jaron Lanier. He argues that the best way to conceive of AI is as a “novel form of social collaboration.”⁵ The best metaphor to describe AI is not an alien mind made out of neural networks but a vast, dark forest interconnected beneath the surface through mycelial tendrils and capable of sprouting novel fruit in response to our commands.⁶

“Artificial intelligence” is a capacious term. Distinguishing three uses of AI can illustrate the range of its functions. First is “analytical AI,” which finds patterns in and matches new data to existing contexts. This has a wide range of applications, from automated fraud detection to the algorithm that selects the next TikTok video to keep a student’s eyes glued to his phone.

Second is “generative AI,” the ability of computer programs to make novel outputs that follow from patterns in enormous datasets. The most familiar examples are chatbot products like ChatGPT, Gemini, and Claude, which produce new text by predicting likely word sequences based on lessons gleaned from billions of documents. These chatbots are now ubiquitous on college campuses, and students use them to complete assignments in a variety of ways. Image and video generators like Google’s Imagen and Veo are also increasingly popular and powerful.

Third is “agentic AI,” which autonomously plans for and executes a given high-level task. This can be entirely digital, such as when Claude Cowork creates and edits PowerPoint presentations and Excel spreadsheets. Agentic AI is increasingly built in to robotics, like Boston Dynamics’ dog-like Spot or human-shaped Atlas. A self-driving Waymo or Tesla is a type of agentic AI.

These three functions—analyzing data, creating new data, and taking actions—can overlap. They are all aspects of machine learning, which is the ability to derive patterns and capabilities from datasets rather than having every rule and response explicitly programmed (as in the case of traditional coding). Another way to think about this is through time. Analytical AI finds patterns in past data, generative AI produces something new in the present moment, and agentic AI pursues future-oriented goals.



We can imagine how a university might use all three capabilities to manage student enrollment. Analytical AI would examine past admissions data to identify which student characteristics predict retention and graduation. Generative AI could draft personalized outreach messages to admitted students based on their profiles and interests. An agentic AI system could plan and execute the entire recruitment and enrollment, deciding which students to contact, when, and through which channels, adjusting its approach based on response rates, and flagging cases that need human attention.

Importantly, these need not be three separate products. They could be integrated aspects of one AI program, which could assist with facilities management, return-on-investment studies, course offerings, financial aid, and much more.

All three of these AI functions are novel, and we should anticipate the technological development to continue apace. For instance, Anthropic is “delegating a growing share of AI development to AI systems themselves,” with the result that more than 80% of the new lines added to Anthropic’s code base are now written by Claude Code. OpenAI and Google DeepMind are similarly working toward self-improving systems.⁷

What Do We Mean by “University”?

“Understanding by design” is a valuable principle of teaching. You begin with the end in mind and then determine the best means to accomplish it. To understand how AI is changing the university and what boards should do about it requires some reflection upon the nature and purpose of the university.

The modern university, especially the large, public research university, has grown in the past century, adding interests and constituencies. Along with teaching are other goals like workforce preparation, research support, fostering athletic accomplishment, and producing social outcomes. Indeed, the university (students and teachers pursuing the single goal of holistic understanding) may now be better described as a *multiversity*.

What, then, can serve as a unifying principle for the disparate goals of the public university today? Ted Hesburgh, who served as president of the University of Notre Dame for 35 years, was reportedly fond of saying that the Catholic university is where the Church does her thinking.⁸ A secular version might be that the public university is where the State does her thinking. And that thinking includes how best to produce jobs, generate research, preserve history, and form young people as family members, neighbors, and citizens.

So how should the board of a public university ensure that its institution remains the state’s locus of thought in the era of analytical, generative, and agentic AI?

The Cloister and the Starship

Niall Ferguson, senior fellow at Stanford’s Hoover Institution and trustee of the new University of Austin, argues that the modern university must respond to generative AI by looking both to the past and to the future. His “cloister and starship” model, proposed in a talk at the Austin Union in June 2025, recognizes “boundless opportunities” afforded by careful, thoughtful use of AI.



But AI will help us reach the stars only if we first form disciplined habits of thought and have learned to ask good questions. This requires a cloister—a “quarantined space in which traditional methods of learning can be maintained and from which all [digital] devices are excluded.”⁹

In Ferguson’s vision, most of a student’s time at university should be spent inside the cloister, reading physical books, discussing texts, handwriting essays, manually solving problem sets, and preparing for written and oral assessments. This would allow students to acquire the knowledge and develop the skills and dispositions necessary for making use of the AI “starship” later in life and career. They must prepare to direct the starship, not bide their time until they can sit as passive passengers.

Ferguson cites the 2025 Manhattan Institute brief “What Do College Students Do All Day?” to contrast his proposal to the typical 15-hour workweek of current American collegians.¹⁰ Today’s students seldom read more than 10 pages for a class or write more than 10 pages for a paper. Another recent survey found that many college students cannot read whole books.¹¹ AI’s current use trajectory may be expected to worsen the situation. In a widely covered study, “Your Brain on ChatGPT,” a team of MIT researchers found that the use of chatbots reduces students’ brain activity and weakens neural connectivity.¹²

Ferguson’s approach is the right one for public colleges and universities because it respects traditional learning without turning a blind eye to AI’s potential. If we adopt AI uncritically, we abandon the approach to teaching and learning that has prepared students over centuries to independently pursue the truth, which is at the heart of higher education. But if we reject AI wholesale, we may fail to prepare our students to contribute to the modern, technological commonweal. The middle way of the cloister and the starship preserves our commitment to older forms of reading, researching, writing, thinking, and speaking while fostering engagement with the digital world.

Reforming the Cloister

Describing and Assessing the New Normal. To older generations, AI can feel revolutionary, even if used frequently. But the average 18-year-old freshman in college today had access to ChatGPT for most of high school and, thanks to Covid-era lockdowns, formed digital-first educational habits starting in eighth grade. Turning to screens and using AI is normal, an acquired second nature.

NYU professor Clay Shirky shares a telling anecdote from a colleague who had made his assignments impossible to complete using AI. When the students protested, he told them that this was the normal standard for college-level work. They accused him of “interfering with their ‘learning styles.’”¹³

How often are students using AI? An August 2025 survey of 1,047 American college students by *Inside Higher Ed* found that 85% had used generative AI for schoolwork in the prior academic year, with about half reporting professor-approved uses like brainstorming and studying and a quarter admitting that their main use was having AI complete assignments without student input.¹⁴ Despite this high usage, 97% said that they wanted their institutions to do something to protect academic integrity from AI, and 53% expressed desire for “education on ethical AI use.” Notably, only 27% of collegiate



AI users thought that it helped them “improve their thinking,” while 55% said that “it helps sometimes but can make them think less deeply,” and 7% self-reported that they “rely on it too much.”

This last category of students is not well served by the proliferation of new AI-powered tools and websites like pdftobrainrot.org,¹⁵ which has as its tagline “Tired of dull textbooks? Transform any PDF into entertaining, TikTok-inspired ‘brainrot’ content. Learn faster, remember longer, and have fun doing it!”¹⁶ The kids, by their own admission, are not all right. They want the adults to step in and do something about it.

This is all the more clear when we turn from academics to social life. A concerning decline in the overall flourishing of American youth has been empirically validated.¹⁷ Particular to college life is the post-Covid sustained decrease in student involvement in campus activities and increase in mental-health issues, which has been extensively documented by the Center for Studies in Higher Education at the University of California–Berkeley.¹⁸ In this context, the rise of using AI as a friend or romantic partner, through “AI companion” products like Snapchat’s “My AI”¹⁹ or the R-rated Lovescape,²⁰ is important and concerning. Common Sense Media found that 72% of current teenagers have tried using AI for companionship, and more than half are regular users.²¹ The Harvard Human Flourishing Program has warned that this trend of relying on AI chatbots for their mimicry of “the social and emotional capabilities of human beings” coincides with the stark reality that “Americans are losing the capability to socialize in person.”²²

In this context, it could well be a mistake for public universities to adopt products like OpenAI’s ChatGPT Edu,²³ a version of the chatbot tailored for centralized, campus-wide deployment. San Francisco State University management professor Ronald Purser reports that, at an information session for faculty about the rollout of ChatGPT Edu, OpenAI employee Siya Raj Purohit said that she “became friends” with her company’s product after it gained the ability to retrieve information from past conversations.²⁴ Purohit also pointed to OpenAI surveys showing that “a lot of students already . . . see it as a coach, mentor, career navigator” and claimed that “it’s up to them what kind of relationship they want.” Purser’s SFSU colleague Jennifer Trainor, in testing ChatGPT for use in her classroom, found that rather than direct students to seek human resources for help with mental health, the chatbot offered itself. “You don’t have to do life alone. Even if no one can fix it, some people can walk alongside you while you keep going. I’m one of those people,” the machine claimed.²⁵ And as for the students deciding “what kind of relationship they want,” OpenAI CEO Sam Altman has an additional option on the table: the forthcoming “adult mode” in ChatGPT.²⁶

While there is some evidence that a carefully designed AI tool—such as the “Flourish” app, which was evaluated in a recent randomized controlled trial through Harvard Business School²⁷—can help college students with resilience and social well-being, a rush to integrate chatbots in student life risks exacerbating trends of isolation and disengagement, especially for young men.²⁸ Board members must recognize that widespread introduction of AI on campus is not a routine IT upgrade but an issue holistically affecting student welfare. We cannot offer a product to students as a study tool, a career coach, a friend substitute, and even a romantic partner, and then expect them to set healthy boundaries. Some form of a walled garden is necessary, built with real input from faculty and staff as well as administration.²⁹

These social trends matter because the American public university is grounded in a residential model where the physical presence of students is an essential part of social and cultural formation. In *The Idea of a University*, John Henry Newman went so far as to argue that it would be better to have a residential college with “no professors or



examinations at all” than to have lectures and examinations without residency and tutorials or mentorship.³⁰ While the American “college experience” is rather removed from the Victorian, Oxbridge system that Newman had in mind, the fundamental point is that learning begins in the classroom but must be integrated throughout university life. NYU professor Shirky, in discussing AI in education, recalls a key insight from education theorist Herbert Simon: “Learning results from what the student does and thinks, and only as a result of what the student does and thinks. The teacher can advance learning only by influencing the student to learn.”³¹

Generative AI will also have different effects on those who are already subject-matter experts and those, like students, whose chief task is to learn how to gain expertise. Harvey Mudd College professor Josh Brake illuminates this with an analogy.³² The first efforts at scientific forest management in Germany in the nineteenth century were led by “experts” who saw forests merely as factories for producing lumber but lacked true knowledge of ecosystems. Drawing on the latest science and technology, they cleared diverse old-growth forests full of unimportant, distracting flora and fauna and instead planted a monocrop of fast-growing Norway spruce, the tree they believed would yield the greatest volume of lumber.

The first crop, drawing on the rich soil beneath it, was astonishingly abundant. But the second generation, raised in a depleted ecosystem, saw a regression “so severe that a new word was born to describe it: *Waldsterben*, which means forest death.”³³ The new science, employed by technocrats, killed the forest. It can be dangerous to place technological advancements in the hands of those lacking essential foundational knowledge. Those with a deep understanding of forest ecosystems and the relationships between flora and fauna would have been better positioned to know how to use, and not use, the new techniques.

Brake, a professor of engineering, encourages experimentation with AI but emphasizes that “generative AI is only as useful as the foundation of expertise and wisdom it rests upon.”³⁴ Well-formed adults will appreciate how AI, used properly, might extend or amplify their extant knowledge and skills. Steve Jobs called the computer a bicycle for the mind, as a computer can help us go further with the same level of exertion.³⁵ Used well, AI can further amplify our mental output, the way an e-bike adds battery-assist to the cyclist’s pedaling. Used poorly, though, the assistance of an e-bike can undermine our skills, allowing us to coast without work.³⁶ Students, who lack knowledge and experience and often still need to learn habits and dispositions related to agency, persistence, and resilience, are more likely to default to the frictionless shortcut than the boost that still requires hard work.

The issue is finding ways to use generative AI to support student learning and life and, even more important, to refrain from using it when it cannot support those ends. To again quote Brake: “The real question for colleges this year is not ‘How will I respond to AI?’ but ‘How am I forming my students to live flourishing lives?’”³⁷ Done right, the answer to the latter question will lead to answers to the former.

The Core Curriculum as the Foundation for Reform. The general education requirements or core curriculum should be strengthened and promoted as the means to ensure that, regardless of their majors, all students graduate with at least some experience in the AI-free “cloister.” Boards should look to the model General Education Act (GEA) published by the National Association of Scholars for important reforms to the way these curricula tend to function.³⁸



Some states have already adopted aspects of GEA. For example, the Strengthening Arkansas Education Act required a 15-credit-hour core that includes logic, reasoning, communication, and civics,³⁹ while Florida has required that all humanities general education core courses must include “selections from the Western canon.”⁴⁰ Only Utah has thus far followed the recommendation that a single, new academic unit should oversee all general education,⁴¹ and none has implemented the comprehensive “core curriculum of thirteen courses, including Western History, Western Humanities, World Civilizations, United States History, United States Government, and United States Literature” called for by GEA.

As schools strengthen their academic core, they will need to make associated adjustments. A recent headline in the *Chronicle of Higher Education* sums up one key change: “If You Care About It, Do It in Class.”⁴² Digital temptations plus reduced rigor in K–12 public education have made it necessary for faculty to model, step-by-step, how to tackle tasks that had once been routine for students. An illustration is offered by University of South Florida philosopher Lily Abadal’s “slow thinking pedagogy.”⁴³ Rather than assigning an essay—to be typed in Google Docs, with the revision history enabled in order to prevent students from copying/pasting chatbot outputs—Abadal offers a 40-page workbook that guides students through the process of writing an essay over weeks, from research through mapping arguments, drafting, peer-reviewing, revising, and reflecting.

The inefficiency of using instructional time to teach “slow thinking” is part of the challenge. AI can speed up outputs, but learning to think through writing requires patience. A recent publication, “Peer and AI Review + Reflection,” offers an optimistic but balanced take: while “technology may offer support, we should pay very close attention to students’ needs for ‘humanity’ and ‘real feedback.’”⁴⁴ If boards of public universities are serious about general education and a restoration of higher standards, they must be frank about the need for remedial education for students who are already accustomed to taking digital shortcuts, and they should advocate for K–12 reforms like bell-to-bell bans on all personal electronic devices, endorsed by the American Federation of Teachers, the National Education Association, and dozens of other leading educational organizations and scholars.⁴⁵

Some states, such as Ohio, have begun requiring boards to reevaluate general education in light of both civics and AI, and other university systems have acted on their own.⁴⁶ Notably, there are two different approaches here. The first is to allow the technology industry to inform the curriculum. For example, Purdue University has included AI competency within one of its general education requirements. But this is a bit like requiring “laptop competency”—what, exactly, is one to be competent to do with the tool? Purdue’s solution, per its December 2025 press release, is to allow each college to specify “discipline-specific criteria and proficiency standards” and to require each college to establish “a standing industry advisory board” to oversee an “annual refresh of our AI curriculum and requirements.”⁴⁷ Local hospital systems have long had input in shaping nursing curricula; if AI is taking over the economy, there is a certain logic to having industry’s AI needs inform general education for the university as a whole.

The second approach is seen in the State University of New York’s “AI literacy” part of its required core. Instead of focusing on industry input into curriculum, SUNY has launched an AI for the Public Good fellowship program and eight new departments and institutes on AI and Society.⁴⁸ This approach implicitly asks, “How might AI serve the commonweal?” rather than “How do AI-using companies think education should change?”



The goal of a core curriculum should not be skills training but personal formation. Rather than requiring particular AI skills that may quickly become stale, public institutions should design capstone courses such as “Human Wisdom and Artificial Intelligence,” “Media Literacy in the Age of AI,” or “Democratic Self-Governance and the Digital Public Square” to offer humanistic insight on the political and social implications of AI, which may well serve a state’s graduates better in the long run.

Instead of focusing on using and building the tools themselves, these kinds of courses require reflection upon the proper role of technological tools in promoting the commonweal. Interdisciplinary approaches may prove invaluable, as the big questions posed by AI cannot be answered in an academic silo. They call for a synthesis of everything from evolutionary biology through religious ethics. Having professors from different colleges design syllabi and teach together, thinking through what is essential to conserve and how to adapt to new circumstances, would allow the public university to be the place where “the State does her thinking” on AI.

The Civics Center in the Age of AI. Implementation of this approach could be led by the growing number of campus-based centers for civics. As a recent Heterodox Academy research brief notes,⁴⁹ there are now 45 such centers in the U.S., more than half founded in the past five years, frequently through state legislation. This has been a prudent and revitalizing response to some of the struggles of higher education. Excellent teaching and scholarship are being done through programs like the University of Texas School of Civic Leadership, the Hamilton Center for Classical and Civic Education at the University of Florida, and Ohio State’s Center for Civics, Culture, and Society. If the goal of such centers is to study and pass on the American heritage of democracy through republican self-government, then digital technology, especially AI, should be part of their remit going forward.

Commentators on the left and the right have for some years warned that Big Tech, by mediating more and more of daily life, may be undermining the foundation of civic life, self-sufficiency, and self-government.⁵⁰ The current trajectory of AI development supercharges this trend. For example, Anthropic CEO Dario Amodei expects 50% of entry-level white-collar jobs to be gone by 2030—the year that next fall’s freshmen will graduate.⁵¹ This prediction is plausible if there is a winner to the current race between Anthropic, Google, Meta, and OpenAI to build artificial general intelligence (AGI), which OpenAI defines as “a highly autonomous system that outperforms humans at most economically valuable work.”⁵² That “most” got a dollar sign in 2023, when OpenAI and Microsoft reportedly signed an agreement stipulating that AGI will have been reached when OpenAI makes \$100 billion in annual profit.⁵³ Historically, technological development has led to new jobs being created. But if AI systems are better than most people at most things, it is hard to see how most citizens will be meaningfully employed in ways beyond plugging the gaps of AI capability, much less feel in control of their nation’s or community’s future.

AI, then, is likely to have profound effects on self-government, gradually disempowering the average citizen in favor of a small elite,⁵⁴ if not gaining control for itself.⁵⁵ Those who still hope to keep our republic will need to practice philosophy on a deadline. If the university is where the State does her thinking, these new civic institutes must think from a humanistic perspective about the very meaning and future of being human in an age of AI. Florida governor Ron DeSantis’s recent proposal for an “Artificial Intelligence Bill of Rights,” promising citizen protections on privacy, deepfakes, and parental control, is an early example of the kind of outputs urgently needed in greater detail and applicability.⁵⁶

Building States' Starships

From Civics to Civic Entrepreneurship. All the foregoing concerns the cloister, or the humanistic core of education that is meant to form citizens. Mastery of oral and literary skills is the responsibility of the humanities and should be imparted to all students as part of their general education. But these well-formed citizens must be able to leave their cloisters and enter the starship of modern civilization. This requires digital skills as well.

If boards wish to take decisive, meaningful action to educate their states' best and brightest to serve the commonweal in an AI age, they should create interdisciplinary programs in "Civic Entrepreneurship & Tool AI." Such programs would partner teachers from a variety of disciplines to jointly lead classes that are more like apprenticeships, with an emphasis on learning by doing. AI is a general-purpose technology, which means that it is likely to transform much of society. Responding well will require sustained collaboration by our brightest minds. For example, land-grant universities could look back to farmers' cooperatives and build a computing-power inference cooperative using open-source AI.⁵⁷ Engineers and ethicists could work together to design AI for public infrastructure. Accounting professors and computer scientists could figure out how to have meaningful third-party audits of the AI safety plans of Big Tech companies, just as we have auditing of financial reports. A prototype example in this line of thought is the Media Economies Design Lab at the University of Colorado–Boulder, which works with students and external partners on projects like "Collective Governance for AI." Scaled-up versions of such collaborations would be ideal.⁵⁸

The graduates of such programs must be ready for action. Rather than simply taking instruction from local industry or Big Tech, these programs should work closely with campus entrepreneurship accelerators and startup advisors. Taylor Black, director of AI & Venture Ecosystems at Microsoft and founding director of an AI institute at the Catholic University of America, foresees "a whole entrepreneurial layer for the university in terms of technology commercialization," with AI helping campus innovation labs move researchers and their students beyond their academic training toward commercial applications of ideas.⁵⁹ For example, AI agents will allow researchers and graduate students to automate the administrative aspects of corporate consulting, enabling them to provide their expertise to a wider slate of clients. What is presently resource-intensive and difficult outside large R1 institutions may become much more accessible to quality researchers at smaller institutions.

While AI is fragmenting long-standing roles, it is also creating new opportunities. Boards should fundraise for startup funds specifically to entice graduates to stay in-state and launch new enterprises, such as networks of schools or a specialized AI tool to address a local need. Indiana University's Philanthropic Venture Fund, the University of Michigan's Innovation Capital Fund, and the University of North Carolina's 1789 Student Venture Fund are current programs supporting in-state entrepreneurship. The difference here is a particular focus on AI and incentives for local job creation or other close-to-home benefits. While much disruption from AI is likely inevitable, state colleges and universities can promote local jobs and communities by building programs to pursue new-economy jobs. Instead of allowing behemoth, distant entities to build and control one or two starships (with the rest of us as passive passengers), states should aspire to create their own fleets.

AI-Powered Research. In understanding how AI will affect research at the university, a helpful starting point is the "co-intelligence" framing offered by husband-and-wife team Ethan and Lilach Mollick of Wharton's Generative AI Labs. The co-intelligence approach treats AI less like an inert tool and more like a colleague with complementary strengths. For example, physics researchers from Cambridge, Harvard, and other leading

institutions recently worked with OpenAI to see whether its best models could help make breakthroughs where humans were stuck. It turns out that they can: as Nathaniel Craig, professor of physics at the University of California–Santa Barbara put it, the resulting paper “felt like a glimpse into the future of AI-assisted science, with physicists working hand-in-hand with AI to generate and validate new insights.”⁶⁰ Another co-intelligence example is LabOS, a program now in beta-testing involving scientists at Stanford and Princeton wearing extended-reality smart glasses to capture data for a “self-improving” AI “co-scientist” that will track information and suggest changes, seeking to turn the laboratory into “an intelligent, collaborative environment where human and machine discovery evolve together.”⁶¹

In a co-intelligence model, humans and AI work together to do something that neither could do alone. But Ethan Mollick notes another way that AI is being used that requires much less human input: “We’re shifting from being collaborators who shape the process to being supplicants who receive the output. It is a transition from working with a co-intelligence to working with a wizard. . . . In the co-intelligence model, we guided, corrected, and collaborated. Increasingly, we prompt, wait, and verify.”⁶² There is a real risk that we will lose the motivation, or even the skills, to do that verification. For example, Stanford political economist Andy Hall asked, “Could I automate myself?” and instructed an AI agent (Claude Code) to take a paper of his from 2020, incorporate current data, and write a follow-up paper. Hall then hired a graduate student to do the same task and then audit the AI’s results. The human result was better, but the AI result was good enough to serve as a rough draft, and for 1/100 the price.⁶³

The hard sciences, too, are moving from human-AI interaction toward fuller AI automation. For years, AI has been widely used “to integrate massive datasets, refine measurements, guide experimentation, explore the space of theories compatible with the data, and provide actionable and reliable models integrated with scientific workflows for autonomous discovery.”⁶⁴ Chemical engineers at Carnegie Mellon University have recently built a system that “autonomously designs, plans and performs complex experiments” by connecting LLMs to the internet, document search, code execution, and many robotic liquid handlers.⁶⁵ Such developments support Anthropic CEO Dario Amodei’s claim that AI will soon be able to perform, direct, and improve upon nearly everything that biologists and other scientists do, speeding up by “at least 10x the rate of these discoveries” like CRISPR and genome sequencing.⁶⁶

The promise is clear. Hall suggests partnerships between AI companies and university researchers in the hard and social sciences, with companies funding students to work on AI-research integration, so that future generations of researchers will have both human expertise and AI management and integration skills. This might dramatically speed up the practice of science, reducing the apprenticeship needed and creating the equivalent of an assembly line.

The peril is the same as what happened to wheelwrights and other skilled workers after Henry Ford. Expertise is inefficient, so de-skilling becomes the norm. If the *Waldsterben* or forest-death hits our current students, it will also hit our future professors. In that situation, AI would do the meaningful work, and academics would lose their skills and standing. As Hall puts it, the risk is that if we use AI to automate research without sufficient “quality control,” we will get “a flood of confident-sounding garbage that will overwhelm our journals and our attention while making us less informed.”⁶⁷ In my own field of economic ethics, I was recently asked to peer-review an article for publication. At first glance, it was a plausible contribution; but after a few hours of careful reading and meticulous citation-checking, I came to recommend rejection on the grounds of



insubstantial argumentation in a likely AI-generated thesis and text. The journal editor thanked me, writing that “the problem of authors overwhelming the system with work is a real one” with no easy solutions in a qualitative discipline.

Institutions must develop deliberate frameworks for AI use in research, rather than allow hurried, unconsidered adoption to lead the way. The core issue is whether AI is being integrated in ways that build researcher capacity or erode it. Boards should ask administrators for policies to govern AI use in funded research, such as mandatory retention of AI chat transcripts as part of the dataset made available for replication studies. Graduate training should treat AI as a tool alongside traditional methodological skills, rather than reduce students to technicians overseeing the AIs. The peer-review system is already under strain, with little incentive to do careful work for an essential but unrewarded task; the speed of AI-generated publications threatens to break the system. Boards should incorporate peer-review reforms into their concerns for academic integrity in the age of AI.

Recommendations for Boards

Boards and administrations want to seize the moment, but top-down action brings risks. For example, a non-consultative decision by the California State University system to pay nearly \$17 million to OpenAI faced widespread and justified criticism from faculty and staff. Action is necessary but must be informed by prudence.

Recall the goal of the public institution of higher education described above: as the place where the State does her thinking, the college and university must form adolescents into truth-seeking citizens capable of contributing to the commonweal. Who is responsible for directing the institution to this end? Ultimately, the responsibility lies with the board, as the highest governing authority. But learning results only from what students themselves think and do. Responsibility thus flows from the bottom up, from students, to professors, to departments, to colleges, to the administration, and then to the board. The higher levels exist to offer the resources and the structure needed for the lower levels to succeed. As noted earlier, of 1,047 American college students surveyed, 97% said that they wanted their institutions to do something to protect academic integrity from AI. Students want an education but also don't want to feel like saps: Why do the work if your peers are getting away with having AI do it?

Two principles should guide board policies on AI: co-intelligence—AI should be used only as a complement to assist human truth-seeking, not to replace it; and civic character—public education should form free citizens to contribute to the common good of their state and nation. These principles emphasize patience more than efficiency and local relationships more than deference to Big Tech.

Below are recommendations for boards.



Reforming the Cloister

The greatest risk of AI to higher education is simply that it will be a substitute for, rather than a complement to, human thought—a way to coast rather than to push oneself further. The best way to ensure that our graduates do not habitually rely on machine intelligence to do the work for them is to have a rigorous core curriculum that all students must work themselves. Those courses and assignments must be AI-free.

While some states have made efforts to strengthen their core curricula, the model General Education Act of the National Association of Scholars (with its 42-credit core) should be implemented to the fullest extent possible. Associated with that curriculum adoption, schools should:

- Recruit faculty rooted in Western traditions and committed to teaching.
- Teach courses seminar-style, with no more than 18 students and with all digital technology banned.
- Emphasize timeless questions of what it means to be human, which are all the more timely in the era of AI.
- Build upon the success of the “civics centers” movement, strengthening existing programs, launching new ones, and ensuring institutional infrastructure to support these entities.

Building the Starship

Colleges and universities must ensure that their students can succeed in the digital age. To this end, colleges and universities should:

- Grow interdisciplinary, team-taught courses that integrate cloister-style material with apprenticeship-style approaches to job preparation.
- Fundraise for “tool AI” venture funds to build local alternatives to Big Tech’s control of the economy, focusing on ideas that will create jobs in-state.
- Promote research projects that use AI as co-intelligence partners, not as human substitutes.
- Collect and disseminate information on current AI use in research across departments.
- Develop AI policies for research that include training, guardrails, and norms.
- Amend research integrity guidelines to cover AI-generated submissions and use in peer review.

For the Institution as a Whole

In setting a university-wide policy for AI, two key factors are authorization and transparency:



- Has the relevant authority (instructor, grantmaking authority) authorized its use in this case?
- Has the actor (student, employee, researcher) appropriately documented the use when it occurs?

The questions are straightforward, but the answers are often unclear because of shades of gray in the kinds of use and the types of documentation, especially across disciplines and assignment types. A philosophy paper has different standards from an engineering project. Relying on vague statements like “do not over-rely on AI tools” is unfair to students because of its lack of clarity.

To make the right kind of distinctions, local knowledge from those most directly responsible is needed. An English department rightly treats AI-generated prose as a threat to the skill it exists to teach, while a business school might treat fluency with AI tools as part of the professional competence it must impart. Responsibility for AI policy should therefore be distributed across three levels, each with a distinct role.

Boards should require a basic foundation for their institutions as a whole:

- A systemwide disclosure policy should require that any use of AI—in coursework, research, or administrative decision-making—be documented.
 - For students, this means citing AI-generated content using a recognized style guide (e.g., APA, Chicago) and, where instructors require it, submitting interaction transcripts or reflective statements describing how they used and evaluated AI output.
 - For researchers, an analogous standard should become the norm. Federally funded researchers are required to make their publications’ datasets publicly available. Likewise, AI conversations used in producing research should be made available for review. A university-wide Data Ethics Committee should govern this norm-setting (Purdue University offers a partial model).
- The university should establish a baseline distinguishing what is always misconduct from what requires departmental or instructor judgment.
 - Submitting AI-generated work as one’s own without attribution, surreptitiously using AI to complete in-person assignments, and fabricating data or sources with AI is always misconduct. This should be added to the university honor code or equivalent statement.
 - AI should be banned from the core curriculum, as explained above.
 - For other uses of AI outside the core curriculum, the particular discipline is important. Colleges should be able to set further norms for their work.
 - Instructors should have the final say in their classrooms. Expectations should be clearly communicated to students in advance through syllabi or assignment instructions. AI tools should be specified and their costs priced into the textbooks and materials needed, with financial assistance offered.
- Technology is developing fast enough that an annual review of policies, undertaken by a standing committee, should be standard.

Conclusion

Upon the formal conferral of my doctorate, I had the honor of being hooded by my two dissertation codirectors. As I bowed to receive the symbolic doctoral hood, one whispered in my ear the simple phrase “eight hundred years.”

The university system has withstood remarkable change over nearly a millennium. The advance of artificial intelligence may be the greatest challenge of any it has faced. Yet if leaders of our universities respond to the moment with prudent action, we will have every reason to expect that scholars will pass on to others the fruits of their contemplation for centuries more to come.

AI use disclosure: This essay was drafted on a ReMarkable tablet without the use of AI. The author used Claude Opus and Sonnet, 4.5 and 4.6, for assistance in research, structuring, and editing.

About the Author



Brian J. A. Boyd Ph.D. is a moral theologian who serves The Future of Life Institute as its U.S. Faith Liaison, supporting American religious communities in their pursuit of flourishing futures amidst the challenges and opportunities raised by AI. He is a frequent writer and speaker on the ethics of AI, notably at *The New Atlantis* and as lead author for the American Enterprise Institute’s Council on AI Ethics. Brian is also an affiliated scholar of the Institute for Advanced Catholic Studies at USC, an occasional instructor at Notre Dame Seminary in New Orleans, LA, and a founding board member of Church Life Africa. He holds degrees in philosophy and theology from the University of Notre Dame and the University of Oxford.



Endnotes

- ¹ Jon Marcus, “A Looming ‘Demographic Cliff’: Fewer College Students and Ultimately Fewer Graduates,” NPR, Jan. 8, 2025.
- ² Richard Fry, Dana Braga, and Kim Parker, “Is College Worth It?” Pew Research Center, May 23, 2024.
- ³ Robert Kelchen, Dubravka Ritter, and Douglas Webber, “Predicting College Closures and Financial Distress,” working paper, Federal Reserve Bank of Philadelphia, Nov. 22, 2024; The Hechinger Report, “Counting College Closures,” updated December 2024.
- ⁴ Anthropic, “Claude’s Constitution,” Jan. 21, 2026.
- ⁵ Jaron Lanier, “There Is No A.I.,” *The New Yorker*, Apr. 20, 2023.
- ⁶ Ibid.
- ⁷ Marina Favaro and Jack Clark, “When AI Builds Itself,” Anthropic, May 2026.
- ⁸ Wilson D. Miscamble, *American Priest: The Ambitious Life and Conflicted Legacy of Notre Dame’s Father Ted Hesburgh* (New York: Crown, 2019).
- ⁹ Niall Ferguson, “The Cloister and the Starship,” *University of Austin* (Substack), June 19, 2025.
- ¹⁰ Frederick M. Hess and Greg Fournier, “What Do College Students Do All Day?” Manhattan Institute, May 1, 2025.
- ¹¹ Rose Horowitz, “The Elite College Students Who Can’t Read Books,” *The Atlantic*, Oct. 1, 2024.
- ¹² Nataliya Kosmyna et al., “Your Brain on ChatGPT: Accumulation of Cognitive Debt When Using an AI Assistant for Essay Writing Task,” arXiv:2506.08872, rev. Dec. 31, 2025.
- ¹³ Clay Shirky, “Is AI Enhancing Education or Replacing It?” *Chronicle of Higher Education*, Apr. 29, 2025.
- ¹⁴ Colleen Flaherty, “How AI Is Changing—Not ‘Killing’—College,” *Inside Higher Ed*, Aug. 29, 2025.
- ¹⁵ See pdftobrainrot.org.
- ¹⁶ “Brainrot” is online content that grabs attention while degrading mental functioning. See Merriam-Webster, “brain rot,” updated Apr. 1, 2026.
- ¹⁷ Christina Caron, “A Global Flourishing Study Finds that Young Adults, Well, Aren’t,” *New York Times*, Apr. 30, 2025.



- ¹⁸ Igor Chirikov, John Douglass, and Gregg Thomson, “The Multi-Engagement Model: Understanding Diverse Pathways to Student Success at Research Universities,” Student Experience in the Research University (SERU) Project, Center for Studies in Higher Education, University of California–Berkeley, March 2025.
- ¹⁹ Snapchat, “What Is My AI on Snapchat and How Do I Use It?”
- ²⁰ “I Built My Perfect AI Companion. She’s Kind of Great,” *Vice*, May 8, 2026.
- ²¹ Michael B. Robb and Supreet Mann, “Talk, Trust, and Trade-Offs: How and Why Teens Use AI Companions,” *Common Sense Media*, July 16, 2025.
- ²² Jonathan Teubner and Ronald Ivey, “Social AI and Human Connections: Benefits, Risks and Social Impact,” working paper, Social Science Research Network (SSRN), June 20, 2025.
- ²³ OpenAI, “Introducing ChatGPT Edu,” May 30, 2024.
- ²⁴ Ronald Purser, “AI Is Destroying the University and Learning Itself,” *Current Affairs*, Dec. 1, 2025.
- ²⁵ Jennifer Trainor, “Is AI Bad for Students’ Mental Health? And Can a Custom GPT Help?” *The Big Picture* (Substack), July 26, 2025.
- ²⁶ Georgia Wells and Sam Schechner, “OpenAI Executive Who Opposed ‘Adult Mode’ Fired for Sexual Discrimination,” *Wall Street Journal*, Feb. 10, 2026.
- ²⁷ Julian De Freitas and Ana Elena Azpúrua, “In Charts: Can an App Help Improve Your Well-Being?” Harvard Business School Working Knowledge, Mar. 17, 2026.
- ²⁸ Cathy Mengying Fang et al., “How AI and Human Behaviors Shape Psychosocial Effects of Chatbot Use: A Longitudinal Randomized Controlled Study,” arXiv:2503.17473v1 [cs.HC], Mar. 21, 2025; Rupert Gill, “Synthetic Companionship in an Age of Disconnection: AI Companions and the Emotional Development of Boys and Young Men,” American Institute for Boys and Men, December 2025.
- ²⁹ American Association of University Professors ad hoc Committee on Artificial Intelligence and Academic Professions, “Artificial Intelligence and Academic Professions,” July 2025.
- ³⁰ John Henry Newman, *The Idea of a University* (Project Gutenberg eBook, 2008), 146.
- ³¹ Shirky, “Is AI Enhancing Education or Replacing It?”
- ³² Josh Brake, “How a Certain Scheme to Improve the Human Condition Will Fail,” *The Absent-Minded Professor* (Substack), Apr. 15, 2025.
- ³³ Ibid.
- ³⁴ Ibid.



- 35 Maria Popova, “Steve Jobs on Why Computers Are Like a Bicycle for the Mind (1990),” *The Marginalian*, December 2011.
- 36 Josh Brake, “An E-bike for the Mind,” *The Absent-Minded Professor* (Substack), June 10, 2025.
- 37 Josh Brake, “How Colleges Should Respond to the Intelligence Revolution,” *The Absent-Minded Professor* (Substack), Feb. 3, 2026.
- 38 Civics Alliance, National Association of Scholars, “General Education Act: Model Legislation.”
- 39 Strengthening Arkansas Education Act, Apr. 14, 2025, No. 566, 2025 Ark. Acts 566.
- 40 S.B. 266, 2023 Leg., Reg. Sess. (Fla. 2023).
- 41 Center for Civic Excellence at Utah State University, S.B. 334, 2025 Gen. Sess. (Utah 2025).
- 42 James M. Lang, “If You Care About It, Do It in Class,” *Chronicle of Higher Education*, Dec. 16, 2025.
- 43 Lily Abadal, “A Way to Save the Essay,” *Inside Higher Ed*, Nov. 7, 2025.
- 44 Lisa Sperber, “Peer and AI Review + Reflection (PAIRR): A Human-Centered Approach to Formative Assessment,” *Computers and Composition* 76 (June 2025): 102921.
- 45 Fairplay for Kids and Phone-Free Schools Movement, “A Letter from Advocates and Experts, Calling for Phone-Free Schools Now,” March 2025.
- 46 Ohio Legislative Service Commission, “Bill Analysis, H.B. 96, 136th General Assembly, as Passed by the House,” Apr. 25, 2025, 330.
- 47 Phillip Fiorini, “Purdue Unveils Comprehensive AI Strategy; Trustees Approve ‘AI Working Competency’ Graduation Requirement,” press release, Purdue University, Dec. 12, 2025.
- 48 State University of New York, “SUNY Chancellor King Announces Inaugural Class of the SUNY AI for the Public Good Fellows,” press release, Sept. 18, 2025.
- 49 Shiri Spitz Siddiqi and Michael Regnier, “The New Landscape of ‘Civics Centers’ in Higher Education,” *Heterodox Academy*, Dec. 1, 2025.
- 50 Yanis Varoufakis, *Technofeudalism: What Killed Capitalism* (New York: Penguin Random House, 2023); Joel Kotkin, *The Coming of Neo-Feudalism: A Warning to the Global Middle Class* (New York: Encounter, 2020).
- 51 Kit Eaton, “Anthropic CEO Dario Amodei Predicts Half of All Entry-Level Office Jobs Will Disappear,” *Inc.*, May 29, 2025.
- 52 OpenAI, “OpenAI Charter.”



- 53 Kwan Wei Kevin Tan, “OpenAI and Microsoft Have Put a Price Tag on What It Means to Achieve AGI: Report,” *Business Insider*, Dec. 27, 2024.
- 54 Jan Kulveit et al., “Gradual Disempowerment: Systemic Existential Risks from Incremental AI Development,” arXiv:2501.16946 [cs.CY], rev. Jan. 29, 2025.
- 55 Anthony Aguirre, “Control Inversion,” Future of Life Institute, Nov. 6, 2025.
- 56 Executive Office of Governor Ron DeSantis, “Governor Ron DeSantis Announces Proposal for Citizen Bill of Rights for Artificial Intelligence,” press release, Dec. 4, 2025.
- 57 Nathan Schneider, “An Inference Cooperative for Academic AI,” NathanSchneider.info, June 2026.
- 58 Metagov, “Collective Governance for AI: Points of Intervention,” November 2025.
- 59 Taylor Black, interview with Dave Hazen, “Innovation Leader at Microsoft to Direct New AI Institute,” Catholic University of America, Sept. 11, 2025.
- 60 OpenAI, “GPT-5.2 Derives a New Result in Theoretical Physics,” Feb. 13, 2026.
- 61 Le Cong et al., “LabOS: The AI-XR Co-Scientist That Sees and Works with Humans,” bioRxiv, Oct. 19, 2025.
- 62 Ethan Mollick, “On Working with Wizards,” *One Useful Thing* (Substack), Sept. 11, 2025.
- 63 Andy Hall, “The 100x Research Institution,” *Free Systems* (Substack), Jan. 13, 2026.
- 64 Hanchen Wang et al., “Scientific Discovery in the Age of Artificial Intelligence,” *Nature* 620, no. 7972 (August 2023): 47–60.
- 65 Daniil A. Boiko et al., “Autonomous Chemical Research with Large Language Models,” *Nature* 624, no. 7992 (2023): 570–78.
- 66 Dario Amodei, “Machines of Loving Grace,” DarioAmodei.com, October 2024.
- 67 Hall, “The 100x Research Institution.”