

Big Data and Artificial Intelligence in Higher Education

Christa Davis Acampora
Christa.d.acampora@emory.edu

This is a prepublication version of a chapter in Academic Ethics Today: Problems, Policies, and Prospects for University Life, edited by Steven M. Cahn; published by Rowman & Littlefield Publishers, 2022

Please cite only the published version:

<https://rowman.com/ISBN/9781538160503/Academic-Ethics-Today-Problems-Policies-and-Prospects-for-University-Life>

College and university leaders have access to an extraordinary array of data they routinely collect for current and prospective students and employees. Increasingly, they are asked to make use of those data for a variety of purposes—many of them linked with core mission functions and responsibilities. Using data to support better outcomes would seem to be an uncontroversial good, though much turns on what constitutes better, which outcomes are given priority, and what trade-offs are required to attain them. Increased pressure to achieve certain performance metrics combined with technological advances that increase efficiency in obtaining, using, and sharing data presents higher education leaders and other stakeholders with distinctive challenges and opportunities in both the near and longer term. This essay scouts just some of that terrain and makes modest proposals for immediate and longer-term actions.

In the context of higher education, at least three general types of data are concerns: 1) those generated in the course of conducting routine campus operations and activities; 2) those generated outside the institution about the institution and any and all of its members (i.e., students, staff, and faculty); and 3) those generated in the course of analyzing and using the

data from either or both sets 1 and 2. Artificial intelligence and machine learning increasingly enable institutions to access, generate, and utilize data of this third type.

Massive amounts of data are generated for nearly every activity on campus and many activities that take place remotely. The university identification card (or digital equivalent through network credentials) tracks holders' campus locations, vehicles driven, buildings entered, meals and other purchases made, materials browsed or accessed, classes and events attended, and more. University-sponsored apps may track precise locations of users within a geofence, on or off campus. IP addresses may disclose where university affiliates are around the world. Learning management systems (LMS) identify users' courses viewed, locations accessed, documents opened, assignments completed, lengths of time engaged with materials. Faculty collaborations are tracked; their funding sources, expenses, and travel monitored; and their social media followed. Recent threats to health and safety, especially perceived threats to economic and national security interests and those arising from the pandemic global health crisis, have accelerated efforts to collect more data and, importantly, to use it, that is, to do more with it.

Applications of artificial intelligence and machine learning can draw on data from the student information system, campus chatbots, learning management systems, website analytics, social media analytics, financial aid records, and the campus ID management system with capabilities of combining those data to identify regular and *irregular* patterns in student activities and behaviors. These can identify levels of student engagement and risks to student success, including student retention. Profiles generated from those data can inform specific

interventions to increase likelihood of success, a seemingly benevolent act, and inform admissions algorithms, which could identify which *types* of students are less likely to succeed and therefore present greater risks to institutional success measures or more significant drains on academic support resources.

Increasingly, faculty data relating to research productivity can predict future academic impact relative to peers and its contribution to the reputation of one's department or academic unit as a whole. While data use agreements might prohibit the use of predictive analytics in actual individual performance evaluations, these data may, nevertheless, influence judgments about major career decisions, including hiring, tenure, and retention, which are prospective judgments about future performance and likely contributions.

As universities outsource various administrative functions and services, as many needed to do when colleges and universities shifted to remote operations during the COVID-19 pandemic, data aggregated and derived from the use of health, wellness, financial, and mental health services became available, accessible, analyzable, and potentially *sharable*. The higher purposes to which these data may be put can be positive. They can provide insights to improve teaching and learning and possibly enhance students' sense of wellbeing. The use of multiple data sources enables university administrators to draw greater connections and, arguably, refine their understanding of specific student needs to support targeted, just-in-time assistance to help students achieve their goals. Chatbots may serve as effective way-finders and improve the student experience more than passive forms of communication such as kiosks, directories, and

web pages. Additionally, data collected from chatbots—student queries and requests—can also inform needed improvements in communication or services, for example, ultimately enhancing the campus stakeholder experience and even opening the possibility of creating a more curated, student-centric experience. These data also augment the demographic data concerning staff, faculty, and students, potentially enhancing diversity on campus. Finally, data that track faculty activity can indicate the returns on investment in research infrastructure and support as well as community impact.

Achieving these gains entails not only the collection and analyses of discrete or combined datasets. The application of artificial intelligence and machine learning (AI/ML) is necessary to achieve predictive insights, determine points of intervention, and make the sorts of assessments of whole fields of inquiry on the order that is required in predicting student performance and faculty impact. Nearly all institutions rely, to varying degrees, on external vendors and equipment in order to make use of big data on their campuses. And nearly all of these data management and analytic platforms require extensive data sharing. This raises significant concerns about the security of these data and the institutional risk and vulnerability institutions face when reliant on them.

The costs of collecting, analyzing, and interpreting these data are high. In addition to the cost of the software platform and any specialized servers it may require, the more data collected, stored, and shared, the greater the need for compliance management, legal services to support analysis and finalization of licensing agreements, and technical services to assess and refine

service contracts. Others have examined the broader costs of the big data era in terms of the vast amounts of energy and human resources required to develop and maintain them.¹ Such costs might include other forms of social and moral capital, including conceptions of our *privacy*, our *personhood*, and, at perhaps the extreme limit, our *humanity*. In what follows, I will explore several of these dimensions before offering some concrete recommendations for university leaders.

Privacy

The topic of privacy is developed elsewhere in this volume and in the rich literature devoted to the theme, including in specialized contexts of education, human resources, and healthcare. Thus, I will highlight only a few key dimensions and offer some cautionary guidance.

A robust legal and regulatory infrastructure guards against breaches of privacy arising from disclosure of personal information, including the Family Educational Rights and Privacy Act (FERPA), which provides for parental access to educational records of their children and limits disclosures to others; the Individuals with Disabilities Education Act (IDEA), which extends greater privacy to the educational records of those it covers; the Health Insurance Portability and Accountability Act (HIPAA), which safeguards health information with emphases on disclosure and consent; and, most recently, the General Data Protection Regulation (GDPR) in the European Union, which has impacted activity globally insofar as it protects its citizens wherever they may interact virtually. Safeguarding personal information is enshrined not only in case law and relevant domains of social life. Most professional associations for university administrators also affirm and integrate professional ethics concerning personal information,

and some specifically highlight duties to ensure data integrity and the importance of protecting personal information. These measures are all oriented to protect individual rights and reasonable expectations of privacy. They do not generally take up protection against the effects of pernicious algorithms, although the GDPR treads into that territory. Nevertheless, it is helpful to briefly examine the conception of privacy that the legal system protects, particularly given that it looms so large.

The conception of privacy that generally informs the legislative protections and professional practices most closely associated with higher education is drawn from ideas developed as extensions of European common law traditions, made more urgent at the dawn of the twentieth century by social and technological advances of an active press and the then new-found ease and speed of photographic technology. Fundamental to this sense of privacy is a right to be left alone². Most modern conceptions of privacy arise from or are linked with our conceptions of autonomy (control over oneself, presentations and representations of aspects of one's life, and control over one's likenesses). The GDPR includes "a right to be forgotten," a right to erasure about information previously collected but which no longer enjoys consent for use or representation. Privacy is also relevant to (if not essential for) other dimensions of a good life and key to achieving what we might describe as personhood.

Privacy and Personhood

These conceptions of privacy affirm the right to develop and enjoy one's individual personality, free of any unnecessary intervention on the part of the state. They also recognize that development of personhood is not simply cumulative—the piling up of one experience after

another—but also iterative, such that it includes taking away information that one might have previously shared but which one no longer wishes to have publicly available. Some argue that control over information about oneself is important for the development of one’s life, as privacy affects one’s relationships with others and even defines and limits the *types* of relationships that one can have.³ This is because our senses of intimacy and closeness are measured and meted, at least in part, by our ability to disclose, withhold, or share information about ourselves. Following this line of thought, breaches of privacy might well inhibit or impair our ability to have certain types of relationships, degrees of separation or closeness. For example, this is one of the bases of rape shield laws and, consequently, some elements of Title IX proceedings, which prohibit evidence of intimate details about the sexual life of the alleged victim and protect against public disclosure of the victim’s identity. In this sense, then, we can see that privacy is connected with our ability to have a say over what is and isn’t shared (by others) about us, to have a measure of control over what representations are permitted about oneself, and to retain certain crucial human possibilities such as initiating, maintaining, and distinguishing a wide range of relationships, including friendship and love. What we share with others and the extent to which we entrust others with deeply personal and detailed information amounts to a type of moral capital, the accumulation and expenditure of which are essential aspects of what makes us human.

AI.Humanity⁴

Much of the legislation devised to protect privacy, as well as popular and academic discussions of it, focuses on protecting individuals. Broader social impacts are relevant, however, and, given the stakes of decision-making in higher education, should be considered in development,

adoption, and deployment of AI/ML systems. Social sustainability, an emergent assessment criterion, might be integrated in future design processes. Concerns for sustainability, generally, include consideration of the management and conservation of resources so that they remain available and accessible for future generations. The broader domain of social goods and their environmental conditions and dependencies are at risk of being overlooked if limited by considerations of the privacy of individuals.⁵ Rapid changes in technological capabilities and developments in artificial intelligence only underscore the need to articulate priorities of social sustainability in the interest of humanity.

The application and integration of artificial intelligence within human social systems is evidence of what some have described as the fourth industrial revolution with an accompanying shift in understanding the place of humans in relation to the rest of the world. This change is characterized by transformations of “systems of production, management, and governance,” driven largely by advances in technologies, particularly those derived from applications of artificial intelligence and machine learning. While some imagine this as a doomsday scenario that radically disrupts the workforce preparation to which higher education contributes, others envision a future of work that enhances and even potentially perfects positive and pro-social human capabilities.⁶

Regardless of the ultimate outcome, there is clear agreement that data—its collection, analysis, and utilization—will drive fundamental changes in virtually every aspect of human life, including how we think about preparation for the world of work, what we think of as research and

discovery, and how we think about teaching and learning.⁷ This is already clearly evident in the Internet of Things (IoT) in which physical objects with imbedded sensors share, access, and process information among and with other objects and systems via the Internet or some other communication system. The objects connected in these ways can also be human bodies on campus.

A company marketing itself as providing “data-as-a-service” has developed stickers called “Bio Buttons” that adhere to the human body for up to ninety days. Their embedded sensors continuously measure and monitor temperature, respiration, heart rate, gait, activity levels, sleep patterns, and bodily positions. The sensors connect with “conversational AI” for remote care. During the pandemic, the company struck a deal with a university seeking community health surveillance and a way to track students on campus, monitor their health for symptoms of COVID-19, and support contact tracing. Insofar as Bio Buttons would provide early alerts for infection and enable tracking even among persons otherwise unknown to each other, the system would know more about certain important features of the lives of all campus constituents than even the individuals themselves could know.⁸

The Internet of Things allows us not only to do more things via connected devices, for example, pay for a cafeteria meal using a smart phone or open the door to one’s home from a remote site, it also collects large amounts of data. Even when such data are depersonalized, they are still collected, aggregated, and potentially indexed to other activities among other devices, and all of this has a value; it can be monetized. These large datasets, rendered intelligible and

actionable with artificial intelligence, can be used to develop profiles, providing information about the interests and inclinations of large groups of people, who may have dozens, even hundreds of distinctive similarities and distinguishable traits that had been unrecognizable by ordinary human intelligence in the past.

Data generated by and among entities in the IoT potentially shed light on users' interests, behaviors, and states. Machine learning can be used not only to identify otherwise indiscernible patterns of similarity but also to devise predictive insights and uses intended to influence behaviors. This has led some to suggest that it is not the IoT that matters so much as the Internet of Behaviors (IoB). The IoT connects physical objects with embedded sensors to exchange information, whereas the IoB extrapolates from those data to understand behaviors and to *influence* them in the future. In a now infamous "experiment," Facebook programmers sought to influence and measure "emotional contagion" in which they manipulated the content of user feeds to test the hypothesis that users could be emotionally directed—both positively or negatively—relative to the emotional content in the newsfeeds of their social networks. This research was conducted entirely without the consent or awareness of the participants.⁹ As a way, in part, to address the widely acknowledged mental health crisis among college-age persons in the U.S., Ellipsis Health developed a machine learning app to identify people with speech patterns matching those with depression. The app, piloted by a college in late 2020, detects changes in tone, pitch, and voice modulation that match the behaviors of those with anxiety and depression. The app is supposed to be driven by an AI/ML engine that enables the program to learn from large data sets rather than rely on specific, predetermined

characteristics. In addition to the partner in higher education, the company is working with a healthcare system to create an app-based visualization of mental health wellness for its clients. While an application like this might be seen as valuable for connecting students with mental health resources before they reach a point of crisis, there is also a risk of pathologizing ordinary human emotions and cultural differences.

Data in the IoB come from a range of sources, including commercial data from transactional purchases, social media and data derived from facial recognition in the vast libraries of personal photos amassed through Instagram and Snapchat, and images captured unwittingly through surveillance technologies.¹⁰ At the current stage of development of AI/ML, personal and personally identifiable information is becoming increasingly less important except insofar as it provides an index connecting other data bits in a profile that can become actionable, that is, a profile that can be used for nudging or, more aggressively, determining the behavior of so-called subjects. Instead of personalized data, there are intersections of datapoints. These may be entirely depersonalized, stripped of all personally identifiable information and perhaps even utterly unrecognizable by the subjects in question, denuded of their agency, so to speak, such that one can refer to *behaviors without subjects*.

This should be worrisome for higher education as one of its higher aims is to enable the development of human capabilities and respect diversity. An Internet of Behaviors could foster homogenization and polarization. If measured, in a primary way, by behaviors, human diversity could be diminished by an IoB that nudges or directs toward specific targets. In this respect, the

concern may be less about what data we are selecting to store and share and more about the ways in which those data arrays create virtual environments that are, essentially, *selecting us*.

Without thoughtful deliberation, not just about protection of personally identifiable information, uses of artificial intelligence on campus could significantly impact the expression of human intelligence in teaching, learning, and discovery.¹¹ Personalized learning environments could limit academic choices, constrain exploration of new ideas, and diminish powers of selection and discernment of learners who might not fit the profiles of those who had succeeded in those fields in the past. In that case, areas of human inquiry could very well stagnate. Chatbots, while optimizing student experiences and interactions on campus, could end up directing behaviors (e.g., funneling some students to certain opportunities and not others) and ultimately shaping their relationships.

Not only are these tools intrusive and subject to malicious tampering, they are also potentially misleading as in a 2021 case of Dartmouth medical school students who were charged with cheating after a large number of students had accessed the learning management system (LMS) during an examination. It was later discovered that the LMS generates usage data even when users are not actively using the system. Students protesting the action underscored the ways in which the incident undermined trust on campus, a key component for creating a learning environment that encourages students to take risks. Remote proctoring services require scanning of surroundings (in most cases during the pandemic, this was students' bedrooms) and monitoring eye and head movements. Persons with cultural backgrounds that differ from

those of the subjects used to train the system and persons with learning differences and disabilities could be flagged as cheating when they are not.¹²

There are similar concerns about automated decision-making in college and university enrollment management practices, including how and where institutions recruit, whom they admit, and what aid they award.¹³ AI applications enable institutions to predict student persistence based on a variety of indicators that have nothing to do with readiness, preparation, or academic achievement and aptitude. Misuse of these tools could adversely affect efforts to enhance student access. For example, colleges and universities can now track prospective student behaviors in their interactions with the institutions' websites, other behaviors on the internet more broadly, as well as interests suggested by their uses of social media. Artificial intelligence and machine learning tools allow institutions to develop profiles comparing prospective students with successful graduates. In a recent university business journal, an admissions officer describes using AI in the applicant selection process, implementing a scoring system that integrates internet behaviors and social media interactions along with data collected from the application itself. In this particular case, accessing the college's library website contributed fewer points to a prospective student's admissions score than pulling up the admissions requirement page because the latter shows greater specific interest in the institution.¹⁴ When interest in the institution itself, an indicator of likely yield (an institutional metric referring to the percentage of students who accept an offer of admission and enroll at the institution) is more valuable than a perceived interest in academics

(presumably a core mission of the institution), then these data are not serving students themselves but rather the institutions recruiting them.

Typical student success indicators, such as first-year retention rates, are not attributes of *students* but rather *institutions*. If an institution applies the metric to the student, then the institutional mission or accomplishment appears to be just being good at picking winners, not actually helping students to succeed. Colleges and universities that recruit and select in this way reproduce a particular type of student rather than identify those most likely to benefit from the educational opportunities the institution can provide. Risks to diversity and institutional development are clearly evident as well. Furthermore, this is also potentially not a winning enrollment strategy in the long run because demographic shifts suggest that, in order to thrive, institutions will need to reach new students and adapt their academic offerings to reflect new fields of knowledge as well as new workforce demands and opportunities.

Recommendations

The potential for bias in the algorithms used in and derived from artificial intelligence is significant,¹⁵ and there is a growing body of research and advocacy for the development of design principles and professional standards that pursue what is sometimes referred to as FATE: Fairness, Accountability, Transparency, and Explainability. Indeed, a growing number of entities are beginning to regulate development and uses of AI, not only with concern for more individual privacy protections but also to minimize or reduce known biases that can arise from the machine-learning and transfer processes,¹⁶ and to protect from discriminatory exercise of automated decisions (e.g., in borrowing and housing) as well as predictive reach, such as in

policing and, as illustrated earlier, in college and university admissions.¹⁷ Leaders in higher education should be supporting the development of FATE standards and accountability measures, integrating them in curricula, and applying them within the campus administrative context, including:

1. Adding literacy standards to include learning objectives related to practices of obtaining information (data acquisition literacy) and standards for learning about how personal information is collected, shared, extrapolated, and applied in other contexts; learning standards should also include understanding how machine-learning works and its potential applications.¹⁸
2. Charging custodians of data with developing data use agreements that ensure standards for ethical use and supporting their integration in statements of professional association ethics. (Various professional organizations affirm privacy values but have not yet integrated value statements connected with responsible uses of data beyond privacy.) Accrediting bodies should expect these to be adopted at the institutional level. These should be specific to heighten awareness and support full community literacy.
3. Ensuring that institutional licenses and service agreements expressly prohibit data brokerage and data retention by third parties, including the creation of datasets and secondary analytical data.
4. Disclosing—in plain language—what data the institution collects and allowing students and employees to opt out of non-essential data collection activities (including a right to be free from nudging).

5. Developing data collection, retention, *and destruction* policies and practices, much like those many institutions have for document collection, retention, and destruction.
6. Nurturing a culture of questioning that can support trust rather than suspicion. Just because we *can* collect data and develop extrapolations doesn't mean that we *should*.

Does *want-to-know* among campus administrators license *should-know*?

Conclusion

This chapter has focused on some significant risks associated with the use of artificial intelligence and machine learning in higher education, highlighting specific ways it is currently implemented. Another risk, however, is that of *not* using it. This risk arises from what would be missed opportunities to reap the benefits of AI to develop novel solutions to large and complex problems (e.g., addressing and mitigating the impacts of systemic and institutionalized racism on generations of students). It is possible that AI could allow us to realize forms of agency and problem-solving on scales that are nearly impossible today. To do this, however, and to avoid the pitfalls that were the focus of much of this essay require more than simply *raising* concerns about ethics. In that domain, too, we are far from a set of agreed upon principles or settled systematic ways of making hard choices that we could build into the algorithmic structure. Fears of uses and abuses of AI could result in holding its applications at arm's length, which could also mean that higher education institutions do not contribute the intellectual resources and commitments to develop best practices identified in the previous recommendations.¹⁹ Thus, efforts to define AI.Humanity should be welcomed, since such a task pushes us to articulate the values that are important in humanity-preserving design principles and critical frameworks.

¹ Kate Crawford, *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence* (New Haven: Yale University Press, 2021).

² Samuel D. Warren and Louis D. Brandeis, "The Right to Privacy," *Harvard Law Review* 4, no. 5. (1890): 193.

³ Charles Fried, "Privacy," *Yale Law Journal* 77:3 (1968).

⁴ I gratefully acknowledge inspiration for theme "AI.Humanity" in the emerging strategic framework of my new colleague at Emory University, Provost Ravi Bellamkonda. At the time of my writing this chapter, the reference was notional, so responsibility for the way I give it flesh here—for better or worse—is my own.

⁵ Wolter Pieters, "Beyond individual-centric privacy: Information technology in social systems," *The Information Society*, 33, no. 5 (2017).

⁶ Klaus Schwab, "The Fourth Industrial Revolution," *Foreign Affairs* (2015). Reprint: <https://www.weforum.org/about/the-fourth-industrial-revolution-by-klaus-schwab>

⁷ Luciano Floridi, *The Fourth Revolution: How the Infosphere is Reshaping Human Reality* (London: Oxford University Press, 2014) 3.

⁸ BioIntelliSense website, accessed September 22, 2021: <https://biointellisense.com/> Oakland University initially required student athletes and dormitory residents to wear the BioButtons; later, it made them optional.

⁹ Adam D. I. Kramer, Jamie E. Guillory, and Jeffrey T. Hancock, "Experimental evidence of massive-scale emotional contagion through social network," *Proceedings of the National Academic of Sciences USA*, June 2, 2014.

¹⁰ Brian Burnke, ed. *Top Strategic Technology Trends for 2021* (Gaertner) p. 4, emphasis mine.

<https://www.gartner.com/en/publications/top-tech-trends-2021>

¹¹ Some largescale potential impacts are charted in Elana Zeide, "The Structural Consequences of Big Data-Driven Education," *Big Data* 5, no. 2: 164-172, doi: 10.1089/big.2016.0061.

¹² See Katherine Mangan, "Dartmouth Dropped a Shaky Cheating Investigation, but Concerns Over Digital Surveillance Remain," *Chronicle of Higher Education*, June 14, 2021. Critiques of uses of AI in proctoring include: 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*, (2020); and Todd Feathers, "Proctorio Is Using Racist Algorithms to Detect Faces," *Vice*, April 8, 2021. These sources are drawn from "EDUCAUSE QuickPoll Results: Artificial Intelligence Use in Higher Education," D. Christopher Brooks, June 11, 2021. <https://er.educause.edu/articles/2021/6/educause-quickpoll-results-artificial-intelligence-use-in-higher-education#fnr8> Accessed August 29, 2021.

¹³ See a Brookings Institute report: Alex Engler, "Enrollment algorithms are contributing to the crises of higher education," September 14, 2021 (<https://www.brookings.edu/research/enrollment-algorithms-are-contributing-to-the-crises-of-higher-education/> accessed September 27, 2021).

¹⁴ Dawn Papandrea, "The Data Pitch: Playing with big-league data to approach and attract students," *University Business* (March 2018), 39.

¹⁵ See Cathy O’Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (New York: Crown Publishers, 2016).

¹⁶ For example, see Joy Buolamwini and Timnit Gebru, "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification," *Proceedings of Machine Learning Research*, vol. 81 (2018).

¹⁷ The GDPR includes a section on profiling. In the main, it is a refusal of automated decision-making in higher stakes areas of life (e.g., credit applications), asserting a right to appeal to a human being. <https://gdpr.eu/Recital-71-Profiling/>

¹⁸ See “AI in My Life’ project,” an effort to develop awareness, understanding, and critical thinking about current uses of artificial intelligence among 15-16-year-olds in Dublin, Ireland. The program is designed specifically for students from underrepresented and lower-resourced communities. A curriculum overview is provided in “AI in My Life: AI, Ethics & Privacy Workshops for 15-16-Year-Olds” <https://dl.acm.org/doi/fullHtml/10.1145/3462741.3466664> DOI: <https://doi.org/10.1145/3462741.3466664> Organizations supporting greater AI education, literacy, ethics, and a more just AI ecosystem include the Algorithmic Justice League, founded by Joy Buolamwi (<https://www.ajl.org/>) and the “AI Now Institute.”

¹⁹ For more on the opportunities and risks, including elaboration of the risks named in this conclusion, see Luciano Floridi, et al, “AI4 People—An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations” *Mind and Machines* 28 (2018): 689-707.