First-Generation, Low-Income Students as Data Subjects in Higher Education Profiling and Prediction AI/ML Applications

Cassidy Pyle†
School of Information
University of Michigan
Ann Arbor, MI, USA
cpyle@umich.edu

Nazanin Andalibi
School of Information
University of Michigan
Ann Arbor, MI, USA
andalibi@umich.edu

ABSTRACT

Artificial intelligence and machine learning applications span myriad contexts ranging from policing to disease diagnostics. While scholars have demonstrated and condemned both potential and extant harms brought about by these technologies, arguably less critical attention has been paid to the ways in which institutions of higher education are leveraging these technological capabilities in ways that may implicate low-resourced college students. We argue that existing AI/ML research articles in the higher education domain sometimes claim to support first-generation, low-income students, but do so without a robust consideration of how their developments may be deployed *at the expense* of these students' self-concept and agency. Furthermore, we assert that these applications produce stigmatizing data bodies around first-generation, low-income students that these students, as data subjects, have little control over.

CCS CONCEPTS

• Human-centered computing • Human-computer interaction

KEYWORDS

Artificial Intelligence, ethics, first-generation college students

ACM Reference format:

Cassidy Pyle and Nazanin Andalibi. 2021. First-Generation, Low-Income Students as Data Subjects in Higher Education Profiling and Prediction AI/ML Applications. In *Human Factors in Computing Systems (CHI '21). ACM, New York, NY, USA, 5 pages.* https://doi.org/10.1145/1234567890

1 Introduction

Artificial intelligence and machine learning tools are deployed broadly in areas such as policing, healthcare, the workplace, and education. In higher education, they are increasingly being used to identify students at risk of academic disengagement and failure [3,12] and to predict retention and graduation rates [7,14]. Research has also been interested in developing and assessing these tools. Research papers explicating these tools often boast about the academic and psychosocial benefits they will bring to disadvantaged students such as first-generation and low-income collegegoers, such as through providing timely feedback and support [4]. Preliminary evidence suggests a dearth of critical attention to ethical issues, specifically potential distributive (i.e., related to material outcomes) and discursive (i.e., related to the meanings that are assigned to groups) harms that may stem from these artificial intelligence tools. In fact, a systematic review notes that a mere 1.4% of reviewed papers detailing AI/ML applications in higher education include a critical reflection upon ethics [13]. A notable exception is Hu & Rangwala's attempt to develop fair machine learning models for at risk student prediction, which moves this corpus of research in a direction that acknowledges that harms may ensue, such as the discouragement of minority students labeled as "at risk" [5]. Still, potential harms of AI/ML applications in higher education for low-resourced students such as first-generation, low-income students deserves systematic in-depth attention. In this position paper, we speculate on profiling and prediction AI's potential harms related to stigma as well as a lack of student agency over their data bodies and chart avenues for future research in this space.

2 AI/ML Applications in Higher Education: A Brief Overview

Broadly, applications of artificial intelligence and machine learning tools in higher education fall into four main categories: profiling/prediction, assessment/evaluation, adaptive systems/personalization, and intelligent tutoring systems [13]. Profiling occurs to determine which students are by and large "at risk" and these profiles are used to predict admission, retention and graduation likelihood as well as academic achievement [13]. Assessment/evaluation concerns feedback, including automated grading [13]. Adaptive systems are generally used in the classroom but are interestingly also used to support student engagement with campus resource centers [13]. Intelligent tutoring systems are less concerned with the administrative elements of higher education and more attuned to what happens in the classroom. These systems teach course content, provide feedback to learners, aggregate learning materials in response to particular student needs, and support collaboration among peers [13].

Based on a review of profiling and prediction AI/ML papers in the higher education context published in 2019 [13], we argue that very few papers in this area explicitly map out the intended uses of built systems, nor do they explicate how exactly these tools translate into positive outcomes for students. Furthermore, they rarely reflect upon ethical considerations and potential discursive and distributive harms. While this valuable systematic review of AI applications in higher education posits the question, "Where are the educators?" [13] we ask, "Where are the students?". In the following section we outline anticipated potential implications of *profiling and prediction tools* that designate some students as "at risk".

3 AI/ML Profiling and Prediction Tools' Implications for First-Generation, Low-Income Students

3.1 Profiling and Stigmatizing Algorithms

That AI/ML is used to profile students as "at risk" renders these tools stigma machines, or "…machines of inscription set in motion through concerted efforts in order to immobilize, wound, humiliate and/or dehumanize those caught within their grasp" [10:260]. Stigma can be levied at first-generation, low-income students in several ways, including from an automated feedback system or from their instructors or peers. Several of the AI applications reviewed involve feedback to the student about their "at risk" status [2], thus demarcating these students and rhetorically separating them from their "normative" peers. Drawing from Feminist Media Studies, HCI scholars have conceptualized algorithmic symbolic annihilation [1], a process by which algorithms perpetuate normative narratives about phenomena and further stigmatize and marginalize people with certain experiences and identities. We argue that this process extends to the case of profiling algorithms in higher education through positioning some students to be the "norm" and in power, and some as "at risk."

Beyond demarcation, the label pathologizes and infantilizes first generation and low-income students without accounting for the structures that endanger them in the first place. It does so in many cases based off of data such as previous grades [2], trapping first-generation, low-income students who routinely struggle to attain a high GPA [6], particularly in their first few semesters, in an endless cycle of stigma. Ultimately, AI tools that engage in profiling "influence, subtly and overtly, how we understand those people" [9:10] which can facilitate stigmatization by associating first-generation, low-income students with risk and implicit assumptions about work ethic and ability.

First-generation and low-income students already carry stigma [11], and we argue that profiling AI/ML tools have the potential to exacerbate both the enacted stigma that students experience (i.e., the microaggressions they may face by teachers, peers, and administrators as a result of being labelled "at risk") and the internalization of this stigma (i.e., they may feel that because a seemingly "objective" tool has labelled them at risk, the stigma that other people assign to them is warranted). These are real harms that should be addressed if we strive for equitable higher education and associated technologies, rather than those that further marginalize and harm.

3.2 Data Bodies and Data Misuse

Data and power operate in a tense and mutually reinforcing relationship. The concept of data bodies reanimates this relationship, arguing that data bodies are "discrete parts of our whole selves that are collected, stored in databases, the cloud, and other spaces of digitally networked flows, and used to make decisions or determinations about us. They are a manifestation of our relationships with our communities and institutions, including institutions of privilege, oppression, and domination." [8:24] This metaphor of data as corporeality raises questions regarding consent and agency. How much control one has over their own data bodies is directly influenced by their social identity and proximity to privilege and power. In the case of low-resourced students and profiling prediction AI/ML in higher education, students have little to no agency over their own data bodies, which are rendered malleable, manipulatable, and analyzable by those in positions of power (e.g., instructors and administrators).

We see in extant profiling and prediction AI/ML applications in higher education the distillation of complex, whole people with rich lived experiences into so-called "at risk" youth at the mercy of instructors and administrators to use this data-based characterization of them equitably and in the best interest of the students themselves. Many first-generation, low-income students both struggle academically in their first few semesters [6] and are lower-income than their continuing-generation counterparts [15], increasing the likelihood that they will rely on aid to pay for college. Their "at risk" status generated by predictive AI can not only influence discursive meaning-making about these students but can also possibly influence administrative decisions regarding who is eligible to access scholarship money or who can appeal academic probation. We go a step further and ask what impact these inferences might have beyond higher education for low-resourced students, for example when graduating and looking for jobs.

The concept of data bodies highlights the ways in which marginalized folks are often denied access to and participation with the very data they create. We argue that those who develop and make decisions to deploy AI/ML tools to be used in higher education have the duty of not only engaging thoughtfully with ethical considerations, but also committing to an ethos of data justice. Data justice is an ethos formulated by Our Data Bodies, an organization that works to understand and combat data-driven inequities that preclude marginalized individuals from thriving. The ethos centers around the assertion that marginalized individuals should have the ability to access and retain common ownership over their data and should be able to equitably participate in conversations and decision-making processes that are driven by this data [16]. What data justice would look like in this context should be determined through systematic engagement with low-resourced students themselves. We speculate that components of data justice in this context would include greater transparency over how student data such as grades as well as identity-based information are being used in these systems. Additionally, it would involve bringing students to the forefront of conversations about how these systems are used, and giving them real power to determine how, when, and to what ends they are implicated in these systems. Transparency should be meaningful, for example it could include opportunities to effectively and freely opt-out of data being used in these systems in a way that does not preclude them from the scholarly opportunities (e.g., timely feedback and support) that these systems may afford. The data justice ethos has the potential to guide technology developers in creating less harmful higher education AI in the future.

4 Conclusion and Future Work

In this paper, we have presented an overview of AI/ML applications in higher education along with a discussion of potential implications for first-generation, low-income students that warrant further research. Potential next steps include: (1) exploratory studies that gauge first-generation, low-income students' awareness of and affect toward AI applications in higher education as well as anticipated harms, (2) design sessions with first-generation, low-income students to speculate on ways in which profiling and prediction AI tools in higher education can be more traceable, verifiable, non-deceptive, and intelligible, which are values noted in [9], and (3) exploratory studies that elicit instructors' and administrators' understanding and use of data bodies about students generated by profiling and prediction AI tools. Those who build AI systems to be deployed in higher education, along with the educators and administrators who play privileged roles in controlling and using the insights these systems generate, should explicitly consider ethical implications of their work as it relates to these student groups and may find the data justice framework useful to do so.

REFERENCES

1. Nazanin Andalibi and Patricia Garcia. 2020. Sensemaking and coping after pregnancy loss: The seeking and disruption of emotional validation online. Proceedings of the ACM Human Computer Interaction, CSCW.
2. David Baneres, Abdulkadir Karadeniz, Ana-Elena Guerrero-Roldán, and M. Elena Rodríguez. 2021. A Predictive System for Supporting At-Risk Students' Identification. In Proceedings of the Future Technologies Conference (FTC) 2020, Volume 1 (Advances in Intelligent Systems and Computing), 891–904. https://doi.org/10.1007/978-3-030-63128-4_67
3. David Bañeres, M. Elena Rodríguez, Ana Elena Guerrero-Roldán, and Abdulkadir Karadeniz. 2020. An Early Warning System to Detect At-Risk Students in Online Higher Education. Applied Sciences 10, 13: 4427. https://doi.org/10.3390/app10134427
4. Ying Cui, Fu Chen, and Ali Shiri. 2020. Scale up predictive models for early detection of at-risk students: a feasibility study. Information and Learning Sciences 121, 3/4: 97–116. https://doi.org/10.1108/ILS-05-2019-0041
5. Qian Hu and Huzefa Rangwala. 2020. Towards Fair Educational Data Mining: A Case Study on Detecting At-risk Students. Proceedings of the 13th International Conference on Educational Data Mining: 7.
6. Jillian Ives and Milagros Castillo-Montoya. 2020. First-Generation College Students as Academic Learners: A Systematic Review. Review of Educational Research 90, 2: 139–178. https://doi.org/10.3102/0034654319899707
7. Zhixin Kang. 2019. Using Machine Learning Algorithms to Predict First-generation College Students' Six-year Graduation: A Case Study. I.J. Information Technology and Computer Science: 8.
8. Tawana Petty, Mariella Saba, Seeta Peña Gangadharan, and Virginia Eubanks. 2018. Our Data Bodies Reclaiming Our Data Interim Report | Detroit. 37.
9. Erica Southgate. Artificial intelligence, ethics, equity and higher education: A 'beginning-of-the-discussion' paper. 20.
10. Imogen Tyler. 2020. Stigma: The Machinery of Inequality. Zed Books Ltd.
11. Deborah M. Warnock and Allison L. Hurst. 2016. 'The poor kids' table': Organizing around an invisible and stigmatized identity in flux. Journal of Diversity in Higher Education 9, 3: 261–276. https://doi.org/10.1037/dhe0000029
12. Jie Yang, Seth DeVore, Dona Hewagallage, Paul Miller, Qing X. Ryan, and John Stewart. 2020. Using machine learning to identify the most at-risk students in physics classes. Physical Review Physics Education Research 16, 2: 020130. https://doi.org/10.1103/PhysRevPhysEducRes.16.020130
13. Olaf Zawacki-Richter, Victoria I. Marín, Melissa Bond, and Franziska Gouverneur. 2019. Systematic review of research on artificial intelligence applications in higher education – where are the educators? International Journal of Educational Technology in Higher Education 16, 1: 39. https://doi.org/10.1186/s41239-019-0171-0
14. Li Zhang and Huzefa Rangwala. 2018. Early Identification of At-Risk Students Using Iterative Logistic Regression. In Artificial Intelligence in Education (Lecture Notes in Computer Science), 613–626. https://doi.org/10.1007/978-3-319-93843-1_45
15. First-Generation Students – PNPI. Retrieved February 8, 2021 from https://pnpi.org/first-generation-students/
16. Our Principles and Values. Retrieved January 22, 2021 from https://www.odbproject.org/about/our-principles-and-values/