Kelly McConvey University of Toronto Toronto, Ontario, Canada kelly.mcconvey@mail.utoronto.ca Shion Guha University of Toronto Toronto, Ontario, Canada shion.guha@utoronto.ca Anastasia Kuzminykh University of Toronto Toronto, Ontario, Canada anastasia.kuzminykh@utoronto.ca

ABSTRACT

The use of algorithms for decision-making in higher education is steadily growing, promising cost-savings to institutions and personalized service for students but also raising ethical challenges around surveillance, fairness, and interpretation of data. To address the lack of systematic understanding of how these algorithms are currently designed, we reviewed an extensive corpus of papers proposing algorithms for decision-making in higher education. We categorized them based on input data, computational method, and target outcome, and then investigated the interrelations of these factors with the application of human-centered lenses: theoretical, participatory, or speculative design. We found that the models are trending towards deep learning, and increased use of student personal data and protected attributes, with the target scope expanding towards automated decisions. However, despite the associated decrease in interpretability and explainability, current development predominantly fails to incorporate human-centered lenses. We discuss the challenges with these trends and advocate for a human-centered approach.

CCS CONCEPTS

• Human-centered computing \rightarrow Empirical studies in collaborative and social computing.

KEYWORDS

Human-Centered Machine Learning, Artificial Intelligence, Literature Review, Higher Education

ACM Reference Format:

Kelly McConvey, Shion Guha, and Anastasia Kuzminykh. 2023. A Human-Centered Review of Algorithms in Decision-Making in Higher Education. In Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23), April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3544548.3580658

1 INTRODUCTION

The use of algorithms for decision-making in higher education, and subsequently the use of student data in algorithms, is growing across the globe [110]. Prior research has found that algorithmic

CHI '23, April 23–28, 2023, Hamburg, Germany

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-9421-5/23/04...\$15.00 https://doi.org/10.1145/3544548.3580658 decision-making has the potential to provide considerable cost savings to higher education institutions, with more personalized and just-in-time service for students [67]. Indeed, students, faculty, and the administration of higher education institutions face growing challenges; increasing tuition fees and debt levels, along with lower levels of government support, have impacted students and higher education institutions alike [110]. Neo-liberalism and the rise of knowledge capitalism within higher education have pushed higher education institutions towards a greater emphasis on metrics, accountability, and KPIs [76]. This has led to the growing use of educational data mining, reliance on learning analytics [1] and the use of algorithms for decision-making, predictions, interventions, and personalization [13].

With this growing trend, both the improper use of student data and the potential for harmful decision-making (predicted and/or automated) within higher education institutions have also risen [67, 107]. In particular, the use of algorithms in education is associated with several ethical challenges, such as student surveillance and privacy, fairness and equity, and interpretation of data [98]. The harmful outcomes of ignoring these and similar ethical issues in algorithm design have recently brought forward significant concerns. The SIGCHI community has considered that algorithm design has failed to identify the true target of intervention [16]. Abebe et al. [2] raise the example of admissions in higher education in that "a computational intervention that aims to equalize offers of college admission across demographic groups might function as a less ambitious substitute for the deeper and more challenging work of improving high school instruction in low-income neighborhoods." Chancellor et al. [25] posit that traditional computational research minimizes individuals to simple data points, and there is cause for concern: machine learning can amplify stigma, reproduce stereotypes, increase discriminatory practices, and harm individuals and communities.

Correspondingly, research is actively exploring different techniques for addressing these ethical issues, including developing frameworks [79], and governance strategies [94], attempting to build more fair models [60, 62, 63, 112], assessing the need for protected attributes as input data [111], and incorporating student perspectives [67, 68].

As the SIGCHI community continues to pursue research on equity and bias in algorithmic decision-making [12, 25, 33, 40, 56, 85, 91, 97], human-centered algorithm design attempts to enable and extend these techniques by incorporating human and social interpretations into the design of algorithmic systems. Specifically, it was suggested [17] that theoretical, participatory, and speculative strategies can be employed to center humans in the design process and to bridge the gap between the algorithm developers and the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

stakeholders who interact with the system or are affected by their decisions. With the significance of the social impact of algorithmic decision-making on various high-stakes human domains [25, 33, 40, 56, 85, 91], there is concern within the HCI research community that the design process of algorithms fails to consider the potential for harm as a result of the inherent uncertainties of the predictions and limitations of the technology [85, 104]. Human-centered algorithm design addresses this concern by leveraging human knowledge from the social sciences and incorporating stakeholder perspectives, allowing researchers to better consider the impacts of algorithmic decision-making in the real world or our context [6].

However, while research in other domains such as child welfare [91], cyberbullying detection [56], law [33], online sexual risk detection [85] and public administration [40] has demonstrated both the lack of and the pressing need for human-centered algorithm design [56], little is known about the application of a human-centered lens in the design of decision-making algorithms in higher education [54, 64].

To address this gap, in this paper, we explore the current trends in the use of computational methods, data types, and target outcomes in the design of algorithms in higher education and analyze the current role and place of human-centered algorithm design approaches in their development. Through a comprehensive review of the existing literature on algorithmic design for higher education (n=62), we collect and qualitatively analyze the models proposed from 2010 to 2021 and demonstrate the existing patterns.

We show that, first, the model design has trended away from rules-based systems towards neural networks and natural language processing. By nature of the model design, the results have become less explainable and less interpretable and increasingly rely on individual student data such as GPAs, enrolment pathways, and Learning Management System (LMS) activity as input features. We also find that the use of protected classes (age, race, gender, disability status) as input features has grown significantly over the past decade, along with the use of the student or applicant's family data such as household income or parental academic achievement (first-in-family). At the same time, while the models are increasingly complex and the decisions become increasingly opaque, the algorithm design does not demonstrate the systematic use of humancentered approaches to reflect the necessary student perspective appropriately.

This work contributes to the community by presenting an indepth account of the current state-of-the-art and trends in algorithmic decision-making in higher education and critically reviewing the algorithms proposed for use in higher education through the application of the human-centered conceptual framework [17]. Moreover, based on our findings, we identify potential gaps in the existing literature and suggest future research opportunities for developing human-centered algorithms for higher education. Building upon existing reviews [56, 85, 91] on the use of humancentered design in algorithm development in different domains, this work provides a foundation for implementing human-centered approaches in the design and development of algorithms in the context of higher education.

In the remainder of this paper, we first review the existing literature on algorithmic decision-making in higher education and the application of Human-Centered Algorithm Design in other

uly 1-26 2022 higher education" + algorithm 9,038 papers identified throug search of ACM Digital Library 5,093 Results 'higher education' + 'predictive analytics 022-07-20 99 results 4,620 duplicate 4 418 titles screened learning analytics 2022-07-25 2,005 results 4005 excluded higher education + machine learning 413 abstracts ot Al in High Educational screened 1,841 results Setting 236 full 177 excluded did not meet retrieved and clusion criter screened lulv 30 2022 62 papers included for synthesi

Figure 1: PRISMA Flow Diagram [69]

domains. We then describe our data collection and data analysis processes, followed by the results of the analysis of the models proposed in the collected literature corpus. Finally, we discuss the critical gaps in the current trends identified through our analysis and propose key opportunities for future research.

2 BACKGROUND

In this section, we provide an overview of the existing literature surveys on the use of algorithms in higher education. We also provide background on the Human-Centered Algorithm Design framework [17] used throughout our research and on how it has been used within other domains. Survey research and systematic literature reviews are important contributions to HCI research, providing insight into what is currently known about the topic at hand, exposing trends and gaps, and identifying opportunities for further research [109]. Previous literature reviews of algorithms in higher education have predominantly focused on model accuracy and performance, rather than fairness through human-centeredness. Nine of the ten reviews that we identified all focused on the same outcome: predicting student success [5, 9, 24, 48, 58, 70, 74, 90, 95]. The remaining study [115] explored various 'AITech' target outcomes. These reviews looked at input data [5, 9, 24, 48, 58, 74, 90, 95], computational methods [5, 9, 48, 58, 74, 90, 95, 115] and model performance [58, 74, 95, 115]. Only one paper [48] included any analysis of ethical issues within its reviewed papers; Hellas et al. [48] briefly raise ethical considerations in their discussion.

Overwhelmingly, the research on algorithms in higher education has been focused on optimizing the algorithms themselves: the inputs, target outcomes, architecture, and performance. The target outcomes aim to make predictions that often impact those humans' lives: what courses they take, what interventions are offered to them, and what programs they are admitted to. Algorithmic performance, though, is a measurement of the functioning of the algorithm and

CHI '23, April 23-28, 2023, Hamburg, Germany

A Human-Centered Review of Algorithms in Decision-Making in Higher Education

Computational Method	Target Outcome	Input Data
Statistical Methods	Grade Prediction	Demographic
Rules-Based	Retention	LMS/Engagement
Machine Learning	Institutional Planning	Institutional
Deep Learning	Pathway Advising	Grade/GPA
Natural Language Processing	Student Services	Enrollment/Pathways
	Admissions	Student Survey
	Assessment	Protected Attribute
	Engagement	

Table 1: Coding Categories

may not align with human and social interpretations of a model's success (does it do and mean what it claims). So while algorithms are developed using staggering amounts of information about humans to then make decisions for humans, it is alarming that humans are so inauspiciously missing from both the design and measurement of their value.

To address these issues, Baumer [17] recommends three strategies for human-centered algorithm design: theoretical, participatory, and speculative.

Theoretical Design, according to Baumer, incorporates behavioral and social science theories into the design of the algorithm. These theories can be used prescriptively to guide algorithm design, informing feature selection, for example, and used descriptively to help us to interpret and evaluate the results of the algorithm.

A *Participatory Design* approach to algorithm design incorporates stakeholders in the design process, the people whom the system will likely impact. In the case of algorithm design and machine learning, this involves connecting the people for whom the algorithm will automate decisions and the end users with the designers, and actively considering the user experience of the system.

Speculative Design, according to Baumer, requires an imaginative approach. Researchers must not only consider the existing circumstances but must extrapolate from it what *could* be. For algorithm design, this requires authors to think through the potential impacts and ramifications of the assumptions and values embedded in the ground truth of the model.

Other domains have demonstrated the feasibility of employing Baumer's human-centered algorithm design framework [17]. In a review of algorithms used in the US Child Welfare system, Saxena et al. [91] found that the literature focused mainly on risk assessment models but does not consider theoretical approaches or stakeholder perspectives. Kim et al. [56] found that incorporating human-centeredness in algorithm design can help develop more practical bullying detection systems that are better designed for the diverse needs and contexts of the stakeholders. The literature [33] also demonstrates how participatory approaches enabled computer scientists and lawyers to co-design Legal AI, and help align computational research and real-world, high-stakes litigation practice. Finally, Razi et al. [85] determined that a human-centered approach was necessary for identifying best practices and potential gaps, as well as setting strategic goals in the area of online sexual risk detection.

3 METHOD

In this section, we describe our scoping criteria and processes for conducting the systematic literature search, coding, and data analysis.

3.1 Literature Corpus

The following keywords were used as search terms to identify relevant papers: "higher education" + algorithm', "learning analytics"', "higher education" + AI', "higher education" + "predictive analytics", and "higher education" + "machine learning". All searches were conducted in the ACM Digital Library between July 1-26, 2022. Our initial search returned 4,418 unique papers. Next, each paper was reviewed for inclusion/exclusion using the following criteria:

- The paper is peer-reviewed published work.
- The paper contains an algorithmic approach and a technical discussion of the computational methods, predictors and outcomes employed.
- The paper uses student data (excluding papers that only used student social media posts) limited to higher education, including university and colleges at any level, and MOOCs (massive open online courses) produced by universities or colleges and aimed at higher ed students at any level.
- The paper was published between 2010 and 2022.

We then cross-referenced the citations of each paper to identify any additional literature that also met our inclusion criteria. We identified sixty-two relevant articles that met our inclusion criteria.

3.1.1 Descriptive Characteristics of the Data Set. All of the papers reviewed (n=62) were published in the ACM full-text collection between January 2012 and July 2022 (Figure 2). The majority (n=59) were published after 2013, with a significant upward trend from 2018 to 2021. We found only one paper meeting the search criteria published before 2012, a scheduling algorithm proposed by Winters in 1971 [108]. It was not included in our review. This is not to say that algorithms were not used in higher education prior to 2012. It is possible that institutions and software providers were developing and using algorithms before this time, or that research was simply published in other venues. The vast majority (n=59) of papers in our corpus were published as conference papers; the remainder were published in journals: Journal of Computing Sciences in Colleges [22, 83], and Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies [105]. The papers came from thirty-five conferences, including the ACM Conference

on Learning @ Scale (n=10) and the International Conference on Learning Analytics & Knowledge (n=17).

3.2 Coding and Data Analysis

3.2.1 Corpus Coding. After identifying the corpus of sixty-two papers, the first author reviewed each paper for its input data sources, target outcome, and computational methods. Our analysis included close reads of the abstract, methodology, and discussion sections for each paper. Papers were reviewed for their specific data input types, target outcome variables, and statistical approaches or machine learning models; coding categories were developed from the results (Table 1). Many papers included multiple models and were therefore included in more than one category of computational method. Similarly, papers also used multiple input data sources and were categorized accordingly. Only one target outcome, however, was identified for each paper. We then conducted a quantitative analysis, determining the number of papers in each category.

Data input type varied greatly between papers. The features include grades, gender, race, first-generation status, prior academic achievement, and enrollment information. This list is non-exhaustive as there was no established standard or pattern for describing input data. In the case of LMS data, for example, it was described by papers as generally as "student interaction data generated in the course" [51] or as specifically as "Time elapsed since last click, Time spent in the course during 7 days, Clicks in time frames, Clicks to date, Clicks in the course during 7 days, Clicks in the forums, etc." [19]. Institutional data includes data specific to the institution (including program or faculty-level data) as opposed to the student: e.g., enrollment and retention rates, geographic data, and financial information such as endowments and tuition rates. Target outcomes were coded as one of the following categories: Student Services, Engagement, Admissions, Grade Prediction, Retention, Pathway Advising, Assessment, or Institutional Planning. Statistical approaches include logistic regression, causal inference, index method, linear multiple regression, Cox Proportional Hazard Regression, Chi-Square test of association, MANOVA, and ANOVA, as well as general references to non-specific statistical techniques. Machine learning models include Arima, Linear Regression, Decision Trees, Logistic Regression, Naïve Bayes, Forward Stepwise Regression, LASSO, Random Forest, XGBoost, Latent Dirichlet Allocation, K-Nearest Neighbors, and various neural nets, amongst others.

In order to analyze the papers in the context of Human-Centered Algorithm Design, we also deductively coded each paper as demonstrating the dimensions of theoretical design, participatory design, and/or speculative design. We adopted Baumer's human-centered algorithmic design framework [17] using the following assessments: *Theoretical Design*: (i) How has the design of algorithms proposed by the papers in our corpus aligned with, or were led by, educational theory? *Participatory Design*: (i) How was meaningful inclusion of stakeholders (students, graduates, applicants, faculty, counselors, and/or administration) realized in the data selection, algorithm design, model evaluated using stakeholder feedback? *Speculative Design*: (i) How have researchers envisioned their proposed algorithms being used in real-world higher-education institutions and scenarios, including the consideration of potential harms and consequences?





3.2.2 Data Analysis. To determine trends, we examined the change in the size of each category over time, from 2012 to 2022. To better understand the relationships between the categories, we also crosstabulated the papers (computational method \leftrightarrow target outcome, target outcome \leftrightarrow input data, and computational method \leftrightarrow input data). Finally, we cross-tabulated our categories with each of the HCAD dimensions: theoretical design, participatory design, and speculative design (Table 7).

4 **RESULTS**

In this section, we present the results of the quantitative analysis of the paper's dataset, structured around the goals of the papers, the predictive features employed, and the computational methods used. We then review the papers for dimensions of theoretical design, participatory design, and speculative design. Finally, we explore the relationships between these elements.

4.1 Input Data

Here we examine the predictors used to develop algorithms in higher education across six dimensions: demographic, learning management system activity (LMS), institutional, grade/GPA, enrollment/pathways, and student surveys. Three papers used for assessment were not categorized as they used assignment-specific input data [35, 64, 89].

Twenty-eight of the sixty-two papers reviewed include student demographic data of some kind in the feature set. The specific input data are shown in Table 2. The most commonly used demographic features were gender, age, prior education, and ethnicity. Two papers [18, 77] use students' personal social media activity (self-reported use of social networking sites and engagement data collected from students' personal pages) for the prediction of grades and retention, respectively. Use of household financial data is increasing, papers included as input features: family assets [30], need for financial assistance [111], and parental income [11, 15, 53, 59].

Five papers have models built only with unique features that could not be categorized. Three are for the purpose of assignment evaluation and use variables specific to the assignment. This includes slide specifications and audio for measuring presentation skills [35, 64], and velocity, acceleration, jerk, and rotation speed in a dental simulator [89]. The remaining two use student behavior (absenteeism, tardiness, uniform violation, and misconduct) [23] and student emails [49] as input data for student services algorithms that provide students with guidance.

In our review, we found almost half (n=27) of the papers include one or more protected attributes (defined as race, sex, age, disability

Year	Total Papers	Gender	Age	Prior Education	Ethnicity	Income	Location	Marital Status	First Gen	Disability Status	Personal Social Media
2012	6	1	1	1	1	1	-	1	-	-	-
2014	6	2	1	1	1	-	-	-	1	-	-
2015	3	1	1	-	-	-	-	1	-	-	-
2016	6	1	1	-	1	1	-	1	1	-	-
2017	15	4	3	3	-	1	2	1	-	1	-
2018	5	2	2	-	-	-	-	-	-	-	1
2019	12	3	2	3	1	-	1	1		1	-
2020	20	5	4	3	3	1	2	1	1	-	-
2021	15	4	3	2	2	2	1	-	-	-	1
2022	2	-	-	-	1	-	-	-	1	-	-

Table 2: Demographic Features by Year

status, or citizenship). Protected attributes, attributes for which discrimination is illegal or protected by some policy or authority, are commonly used as algorithmic inputs in higher education [111]. Some of these attributes, minority status, and family income, for example, are substantially correlated with higher education dropout rates [32]. However, Yu et al. [111] found that the inclusion of protected attributes did not improve their model's performance for dropout prediction, and the use of these attributes as input data may amplify inequities already existing in higher education institutions [112].

Thirty-two papers include student grades as a predictor. This data encompasses course grades, assignment grades, GPA, transcripts, and standardized test results. Student enrollment is also used frequently (n=18) in the form of input data related to course load, course descriptions, educational pathways, and currently enrolled courses.

Authors frequently use interactions with the institution's learning management system as measures of student learning behaviors and student engagement in their respective courses. Eighteen of the papers reviewed used LMS data as input data for their algorithms. Examples of LMS data include timestamps and counts of clicks, time spent watching videos or reading course content, use of discussion boards, and login frequency.

Only six of the papers reviewed include student survey data. The surveys include perceptions of the LMS [66], self-reported mental health data [105], skills, interests, and preferences [73], student services needs [49], self-reported learning styles [68] and prior knowledge assessment [36]. Finally, few of the reviewed papers (n=3) had input data related to the institution or course more generally, and all three were published in 2021. These features included enrollment and retention [10, 83, 118], and digital competency of the institution and faculty [10].

The feature selection process is an important element of model design. Guyon et al. [44] identify three main objectives of variable selection: improved model performance, improved model efficiency (both speed and cost in acquiring and cleaning data, and in running the model), and model explainability. The papers in our review overwhelmingly prioritized model performance and data availability. Explainability was not detailed as part of the feature selection process nor was the process grounded in educational theory. A human-centered approach to algorithm design should begin with a contextual understanding of the data [27] as the curation of input data creates opportunities for bias, algorithmic harm, and privacy concerns. The features selected become the 'ground truth' of the model but the feature selection process is inherently subject to interpretation [27]. In the majority of the papers in our review, the roles of gender, age, and ethnicity became the 'ground truth' in predicting grades or student success. The apparent correlation between these personal attributes and a student's academic achievements is unquestioned, and few papers adopted a human-centered lens to the selection process. Missing from the papers is evidence of an evaluation of the features through educational or social science theory, to speculate on the possible impact of using these features in decision-making, or to include stakeholders in the feature selection process.

4.2 Computational Methods

In this section, we discuss the computational methods used to develop algorithms. For the purpose of analysis, methods and model types were categorized as Inferential Statistics, Rules Based, Machine Learning (ML), Deep Learning (DL), and Natural Language Processing (NLP).

The use of machine learning algorithms within higher education grew rapidly beginning in 2014, as shown in Table 3. As expected, inferential statistics were primarily used prior to this trend. As machine learning methods advanced, researchers moved away from statistical methods.

The overwhelming majority of papers in our corpus, 48 (77%), used some form of machine learning. Almost half of the papers (n=29) used more than one type of machine learning algorithm (46% of all papers, and 60% of machine learning papers). These papers have a stated goal of comparing models for the task at hand.

From 2018 to present, NLP and deep learning both saw rapidly increasing use, indicating a shift towards less explainable and less interpretable models and results. Only four of these papers included any discussion of interpretability and explainability, including feature importance. To increase the interpretability of recurrent neural networks and determine feature importance, authors used the permutation feature importance algorithm [105], random forests [8] and SHAP global feature importance [14]. While these methods

Table 3: Computational Method by Year

Year	Total Papers	Inferential Statistics	Rules Based	Machine Learning	Deep Learning	Natural Language Processing
2012	2	1	-	1	-	-
2013	1	1	-	-	-	-
2014	7	3	-	5	-	-
2015	1	-	-	1	-	-
2016	3	2	1	-	-	-
2017	5	2	-	3	-	-
2018	8	-	1	6	2	1
2019	9	2	-	5	4	1
2020	11	1	-	9	3	1
2021	11	1	-	8	3	1
2022	4	-	-	3	2	1

offer insight into how the algorithm's decision was weighted, only one [14] discussed potential interpretations of the results.

Computational methods can have a meaningful impact on the potential for algorithmic harm. In the case of machine learning, model selection affects more than just the algorithmic performance; explainability and interpretability are key factors ensuring algorithmic decision-making is as fair as possible and human-centered [59]. As Rudin [88] noted, there remains a pervasive myth in the research community that the model selection process inevitably includes a trade-off between model accuracy and interpretability, "a widespread belief that more complex models are more accurate, meaning that a complicated black box is necessary for top predictive performance." We saw this belief shape research in many of the papers we reviewed; with model performance as the primary metric and goal, explainability and interpretability are dismissed in favor of deep learning methods. In order to develop human-centered algorithms for higher education, students must be centered in the model selection process by prioritizing results that can be understood and explained.

4.3 Target Outcomes

In this section, we examine the target outcomes of algorithms proposed in the papers. The papers were sorted into nine categories based on the goal of the model and the target variable: grade prediction, retention, institutional planning, pathway advising, student services, admissions, assessment, and engagement. Table 4 depicts the distribution of papers across target outcome categories, by year.

Thirty-nine (63%) of the papers in our review propose algorithms that seek to predict student success, defined as retention or dropout prediction (n=19), grade prediction (n=18), and admissions decision-making (n=2). While still a significant focus of the research community, student success papers have accounted for a smaller share of the research in recent years.

Student services (n=4) and pathway advising (n=7) make up 18% of the papers. The increase in available data from learning management systems and self-serve course registration may account for the recent increase in research in these two areas.

The shift in focus from student success prediction for retention towards pathways advising and admissions is significant. Retention and drop-out prediction models are proposed to provide insight to academic staff by identifying students in need who may not have been otherwise supported. Pathway advising and, to an even greater degree, admissions models have the power to make decisions for students, acting as gatekeepers to courses and programs.

Seven papers focused on course-level predictions, including predicting student engagement (n=4) and student assessment (such as automated assignment evaluation) (n=3). This is another growing area of research that is also dependent on the influx of LMS data. The remaining five papers related to institutional planning. There is growing variability in the goals of the models proposed by the research community. Ten years ago, models were largely limited to grade prediction and retention. In the last few years, there has been a trend towards student service models: models that students interact with directly for information and advice, such as pathway advising, administration services, and learning support.

4.4 Human-Centered Algorithm Design Strategies

The vast majority of the literature in our review does not include educational, social, or behavioral science theories in the design of their algorithms. Only two of the papers ground their algorithmic design in established learning theory: Borella et al. [19] consider individualistic and constructivist framings while designing interventions for their drop-out prevention tool, and interest exploration, which the authors consider a fundamental component of constructivist learning, is critical to the design of the educational pathways tool developed by Chen et al. [26]. All sixty-two papers in our corpus focus on students and learning, but despite the subject matter, only these two papers make substantial reference to educational or learning theory.

Despite the participatory nature of higher education, only three papers in our review [26, 39, 81] include a participatory approach to their model development. All three have pathway advising as their targeted outcome. The participatory strategies employed include a formative study on need analysis [26], post-intervention evaluation studies [26, 81], and a robust participatory action research approach directly involving academic advisors [39].

While no papers demonstrate a robust speculative design, three papers [53, 59, 111] do consider the use of the tool in serving realworld purposes. For all three, that consideration includes warnings of potential consequences and the "broader implications for using predictive analytics in higher education" [111]. One paper suggests real-world mitigation techniques by advocating for human-in-theloop implementation [59].

4.5 Relationship between Input Data, Methods, Outcomes, and HCAD Strategies

In this section, we examine the trends in the interactions between the input data, computational methods, and target outcomes. We then discuss the relationship between each of the model parameters and each of the human-centered algorithm design strategies: theoretical design, participatory design, and speculative design.

Relationship between Input Data and Target Outcomes

Table 5 crosstabs between the input data used and the target outcomes they seek to predict. GPA/Grades are used as a predictor for all target outcomes at least once, with the sole exception of algorithms developed to assess individual assignments. Demographic data, which includes the protected attributes, is the next

Year	Total Papers	Grade Prediction	Retention	Institutional Planning	Pathway Advising	Student Services	Admissions	Assessment	Engagement
2012	2	1	1	-	-	-	-	-	-
2013	1	1	-	-	-	-	-	-	-
2014	7	3	2	-	-	-	-	2	-
2015	1	-	1	-	-	-	-	-	-
2016	3	1	1	-	-	-	-	-	1
2017	5	2	3	-	-	-	-	-	-
2018	8	2	1	1	1	2	-	-	1
2019	9	2	3	-	3	-	-	-	1
2020	11	3	4	-	1	1	2	-	-
2021	11	3	2	4	1	-	-	-	1
2022	4	-	1	-	1	1	-	1	-

Table 4: Target Outcome by Year

most commonly employed predictor and weighs heavily for grade prediction and retention. LMS activity is applied not only in predicting engagement but also for grade prediction, pathway advising, and retention. Two papers use student surveys for the purpose of grade prediction, a self-report Learning Style Inventory survey [20] and a prior knowledge survey [36].

Relationship between Computational Method and Target Outcome

The computational method used and the correlating target outcome are also shown in Table 5. The review revealed a high variety of computational methods across target outcomes, except in the case of admission prediction (machine learning only). Within those two papers [7, 100] however, a variety of machine learning algorithms are employed. Papers with the most common target outcomes, Grade Prediction, and Retention, include statistical, machine learning, and deep learning methods. Natural Language Processing, which is only found in papers published in 2018 or later, is used only for Institutional Planning, Pathway Advising, and Student Services.

Relationship between Input Data and Computational Method

Next, we cross-examine input data by computational method (Table 6). Papers employing statistical methods and machine learning techniques both used a wide variety of features as predictors. This is to be expected of machine learning, a modeling technique that requires heavily engineered feature sets and, within the papers reviewed here, is used to predict a very wide variety of target outcomes. Statistical methods, however, are less expected as the papers using this technique are limited to four target outcomes. For Grade Prediction, papers employing statistical methods use input data across five categories. For the same target outcome, papers employing deep learning methods use only demographic, GPA, and, for a singular paper, enrollment data. The most recent technique, Natural Language Processing, is associated with only three feature categories: GPA (1 paper [28]), student survey (1 paper [49]), and enrollment/pathway data (4 papers all using course descriptions [26, 28, 80, 81]).

Relationship between HCAD Strategy and Predictor, Target Outcome and Computational Method

In addition to the model parameters described previously, our review also considered the design of the study through the lens of Baumer's [17] framework described in section 2. In this section, we'll identify which papers employ theoretical, participatory, or speculative dimensions in their design, as demonstrated in Table 7. In the next section, 5, we'll discuss the impacts of this inclusion on the literature.

Only one paper [26] included more than one element of Baumer's HCAD framework [17] (theoretical and participatory design). No papers to our knowledge include all three dimensions. All feature categories except *Institutional* and *Student Survey* were accounted for in papers using HCAD strategies. Interestingly, while three papers [26, 39, 81] using participatory design included participant surveys in their design methodology, they did not include surveys as feature sets. Only one paper [49] used student survey data in its feature set, in the design of a student services chatbot. The HCAD papers generally targeted the more common outcomes of *Grade Prediction* and *Retention*, though two papers [26, 81] were focused on *Pathway Advising*. All of the HCAD papers employed machine learning techniques including machine learning (4 papers), deep learning (2 papers), and NLP (3 papers).

5 DISCUSSION

5.1 Research Challenges and Opportunities for Algorithmic Decision-Making in Higher Education

Three major needs emerged from the trends identified by our review: the need to establish valid and theory-informed ground truths, the need to consider interventions within the algorithm design process, and the need for governance as algorithms move from identifying and informing to decision-making.

5.1.1 Establishing Theoretical Groundedness. We found a growing trend toward the use of Learning Management System (LMS) data as input features in algorithm design (n=18). These models were built with the goals of predicting grades (n=8), retention (n=6), engagement (n=3), and pathway advising (n=1). In each of these papers, student interaction with their web-based learning platform was used as a proxy for their engagement in the course. The underlying

		Input Data					Сот	nput	ationa	al Me	thod
Target Outcomes	DM	LMS	IN	GPA	ENR	SS	SM	RB	ML	DL	NLP
Admissions	1	-	-	1	-	-	-	-	2	-	-
Assessment	-	-	-	-	-	-	-	-	3	1	-
Engagement	1	3	-	2	1	-	-	1	3	-	-
Grade Prediction	10	8	-	13	4	2	6	-	11	6	-
Institutional Planning	-	-	3	2	3	-	1	1	3	1	1
Pathway Advising	-	1	-	2	5	1	1	-	3	2	3
Retention	15	6	-	11	5	1	4	-	14	2	-
Student Services	1	-	-	1	-	2	-	-	2	2	1
	DM: Demographics					SM: Statistical Methods				ds	
	LMS: LMS/Engagement					RB: Rules-Based					

Table 5: Input Data and Computational Method by Target Outcome

IN: Institutional GPA: Grade/GPA **ENR:** Enrollment/Pathways

SS: Student Survey

ML: Machine Learning **DL:** Deep Learning NLP: Natural Language Processing

assumption is that LMS activity like clicks on videos, frequency of logins, and views of course content is correlated with their interest and participation in the material, yet in none of the papers was this assumption questioned or was educational theory presented to support it. LMS interaction as a proxy for engagement, and an indicator of academic success, is then accepted as a ground truth within the algorithms. LMS activity, though, can take many forms and there was no discussion within the papers as to the predictive power of specific types of activities. The viewing of a video may not have equal engagement value as participation in a forum discussion, or infrequent but substantial logins may or may not be as effective as frequent, brief ones. Additionally, activity can be impacted by many factors, including technological factors such as interface design and privacy support [114] and the combination of the LMS itself, the course design, and the course type produce different student experiences and influence student behaviors [34]. These factors could differ significantly from course to course, even within the same program or institution. The algorithms in question don't account for the course design. The ground truth of these models is that students with less LMS activity are less likely to succeed, the output of which will trigger a student-level intervention and not a review of the LMS, course design, or user experience.

Beyond incorporating educational theory in the design of algorithms, computational scientists should allow for more collaboration with education subject matter experts and social scientists, to analyze the importance of these LMS interactions critically. More research is needed to understand the relationship between LMS activity and engagement. Additionally, researchers must consider how the importance of the individual features of their algorithms aligns not just with target variables, but also with the greater goals and outcome of the algorithm. We discuss the importance of proposing appropriate interventions as part of the algorithm design process further in the next section.

5.1.2 Considering Interventions within Algorithm Design. We found that before 2018, research focused primarily on designing algorithms to predict students' grades and retention. The motivation behind this research is to reduce attrition through the early detection of students who are identified as likely to drop out or fail a course or program. Once identified, higher education institutions can provide these students with timely and focused interventions. Papers frequently cite significant consequences to students, academic staff, and higher education institutions as a result of attrition, but few consider what interventions exist, the efficacy of the interventions, or the institutions' legal and ethical obligations to provide interventions once 'at-risk' students have been identified [84]. Intervention efficacy, as well, is difficult to measure - interventions varied in form and included automated email notifications, supplementary readings and assignments, recurring meetings with advisors, and counseling services [47]. The one paper identified that did consider the impact of these interventions [31] found no evidence of a subsequent effect on retention outcomes. While dropout prediction was and remains a popular avenue of research, we did not find any literature demonstrating that the models are put into use effectively within higher education institutions. The algorithms seem to be designed within data-science silos, with little consideration for their role within the institution or as part of the greater student experience. A participatory design approach that incorporates learning strategists, for example, allows for the iterative design of theoretically grounded interventions in tandem with the development of the algorithm; any limitations and potentially harmful impact of the model (such as false positives) can be explained and considered in the design of the interventions [21].

While more research is certainly needed to determine if algorithms developed to identify 'at-risk' students and subsequently trigger early interventions have any impact on retention, the acute risk to students appears minimal: at worst the resulting outcome is participation in an ineffective intervention. Student support services at higher education institutions can be scarce, however. Support that goes to a student flagged algorithmically may be support that would have otherwise been provided to another student. In

Table 6: Input Data by Computational Method

	Input Data						
Computational Method	DM	LMS	IN	GPA	ENR	SS	
Statistical Methods	5	7	1	6	2	3	
Rules-Based	-	1	-	1	-	-	
Machine Learning	21	9	2	21	12	2	
Deep Learning	6	2	-	9	5	1	
Natural Language Processing	-	-	-	1	4	1	

DM: Demographics LMS: LMS/Engagement IN: Institutional GPA: Grade/GPA ENR: Enrollment/Pathways SS: Student Survey

other words, within higher education institutions, interventions may be a finite resource. We were unable to find any literature that compared the outcomes for students identified by an algorithm to those identified through traditional methods.

To address these gaps, future research must include consideration of the interventions proposed as an outcome of the model. An algorithm designed to increase retention cannot only be evaluated by its performance in predicting which students may not succeed. In other words, researchers should consider the relationship between the specific predictors used in making the prediction and the intervention proposed to then correct the prediction.

5.1.3 Moving from Identification to Decision-Making. Even more concerning, however, is the move from retention and grade prediction toward pathway advising and admissions. The decisions made by retention algorithms create access to support services that may not have otherwise been available to students but they do not restrict students' choices. Algorithms designed for course selection, program admission, and pathway advising have the potential for a more direct and limiting impact on students. For example, researchers [100] used the same algorithmic target variable, retention, along with likelihood of job placement, to develop an admissions algorithm. In this case, retention was not used to flag enrolled students in need of assistance, but rather to deny applicants admission to a Master's program. In another example, researchers [52] used grade history to build a course recommendation engine, guiding students in what courses to pursue based on likelihood of success. In both these cases, grade prediction is used as a tool to restrict students' options. Algorithms in higher education are moving beyond the ability to simply provide greater insight to institutions through the identification of at-risk students, but to actually make the decisions for them.

As we expect to see this trend continue, and the use of algorithms for decision-making expand beyond its current scope, our results suggest the need for more governance within higher education institutions. Our review indicates that it is certainly technologically possible to leave these student decisions in the hands of algorithms, but whether it is desirable for algorithms to shape students' academic careers without the oversight of a human-in-the-loop remains a legal, ethical, and pedagogical question.

5.2 Towards Human-Centeredness in Algorithmic Decision-Making in Higher Education

Without the application of a human-centered lens, much of the discussion above would have been missed. Embedding Baumer's framework [17] enabled us to center students within our review, and subsequently identify the limitations of the papers themselves and the increased potential for harm within the changing trends of algorithm design.

Higher education institutions are socially-complex, diverse, and high-stakes environments with a potentially vulnerable population of students and inherent power imbalance. Algorithmic systems designed to provide insights into, or more recently, to make decisions for, the student population are inherently part of these systems. Complex and highly contextual student data is used to train algorithms that are then used to make complex decisions for those students. The research community has raised concerns about the use of AI that includes machine autonomy, the consequences of which are quickly being realized [96]. The AI Incident Database strives to document these risks and impacts, and includes many examples of algorithmic harm in higher education including facial-recognition software for exam proctoring providing "allegedly discriminatory experiences for BIPOC students" [46] and wrongfully accusing students of academic dishonesty, an application screening algorithm that "allegedly exacerbated existing inequality for marginalized applicants" [45], and a grading algorithm that "kept students out of college" [101].

A human-centered approach is required to ensure the algorithm is designed with an understanding of those contexts and the realworld functioning of algorithmic decisions must be grounded in social science theories. By including social science theory in the design of their algorithms, researchers can create more robust feature selection processes. Educational theory allows us to more efficiently narrow down potential features to only those that are pedagogically sound, avoid bias, and align to outcomes. Educational theories will also provide a descriptive framework, improving the interpretation of input data and outcomes. Additionally, they provide new avenues for measuring model results beyond model performance. Theoretical design is by nature multi-disciplinary, requiring the expertise of multiple and varied domains. In practical terms, this means that researchers must turn to behavioral and educational theory to understand how their stakeholders work and learn, and use that understanding to select their feature sets, interpret the results, and recommend the corresponding algorithms.

Participatory design is frequently proposed within critiques of machine learning as a strategy for identifying and mitigating risks [33], within the public service[92] and other workplaces [41]. Developing an algorithm to make decisions for students requires the formalization of that decision. As such, model developers must take care to ensure that formalization reflects the true goals of all stakeholders [2]. Abebe et al. see this as an opportunity for institutional reflection, offering the opportunity for non-technical stakeholders to explicitly consider how decisions should be made [2]. By including domain stakeholders, researchers can develop systems that are more fit for real-world use in socially complex contexts, such as those in higher education institutions. Moving towards

		HCAD Strategy					
		Theoretical	Participatory	Speculative			
Target	Admissions	-	-	-			
Outcomes	Assessment	-	-	-			
	Engagement	-	-	-			
	Grade Prediction	-	-	2 [53, 59]			
	Institutional Planning	-	-	-			
	Pathway Advising	1 [26]	2 [26, 81]	-			
	Retention	1 [19]	1 [39]	1 [111]			
	Student Services	-	-	-			
Input Data	Demographic	-	-	3 [53, 59, 111]			
	LMS/Engagement	1 [19]	-	1 [59]			
	Institutional	-	-	-			
	Grade/GPA	1 [19]	1 [39]	3 [53, 59, 111]			
	Enrollment/Pathways	2 [19, 26]	3 [26, 39, 81]	2 [53, 111]			
	Student Survey	-	-	-			
Computational	Statistical Methods	-	-	-			
Method	Rules-Based	-	-	-			
	machine Learning	1 [19]	1 [39]	2 [59, 111]			
	Deep Learning	-	1 [81]	1 [53]			
	Natural Language Processing	1 [26]	2 [26, 81]	-			

Table 7: Human Centered Algorithm Design Strategy by Target Outcome, Input Data, and Computational Method

a participatory design approach will allow for the active involvement of students, faculty, counselors, and administration in the model design process. Engaging these stakeholders throughout the design process will enable algorithms to be designed to integrate properly into existing systems. Including staff of higher education institutions, such as counselors and retention teams, in the design of drop-out predictions algorithm will ensure the model results and features of high importance align with the available interventions. It will also allow for a better understanding of input data, thus improving the interpretation of model results.

The majority of papers in our review evaluated the models on their performance, usually their accuracy, and didn't consider if they delivered on their purpose. Beyond design and development, a participatory approach to model evaluation enables researchers to go beyond performance metrics, evaluating models not just on their technical accuracy but on their ability to meet the goals of the research. The algorithms may accurately identify the students it has been trained are 'at risk', but whether that identification led to meaningful interventions that subsequently increased retention is unknown. Many questions around model value, as opposed to model performance, are unanswered: do drop-out prediction models improve retention; does LMS data reliably correlate to the student experience; are students satisfied with the courses selected for them through a pathway advising tool? Without the perspectives of a participatory approach, future research will be limited to defining success by an algorithm's ability to predict that which it was trained to predict. This is, of course, a crucial step in model development. But to reduce the potential for algorithmic harm, researchers must all consider the success of their models in the context of their intended complex systems. This evaluation requires the participation of stakeholders both during the design of the model and after it is implemented.

Every algorithm in our review was proposed for some real-world purpose, from triggering interventions for at-risk students and reducing attrition to grading assignments. But the literature largely fails to consider the algorithms as part of larger systems, and how they will affect those systems going forward. The research included in our review considers model performance only, with questions of how the decisions will be governed and implemented unanswered, and little speculation about how automated decisions could shape student experiences and outcomes. Future research into algorithmic decision-making in higher education must go beyond model performance to consider the potential the algorithms could have once used for their proposed real-world purpose. To design algorithms that have the intended impact, researchers must consider both the possibility of harm caused by the model but also what changes, if any, are necessary to ensure the algorithm is able to help meet institutional goals.

5.3 Limitations and Future Work

Our systematic literature review was limited to papers published in the ACM Digital Library between 2010 and 2022. We may have missed algorithms published outside this time period, as well as in other libraries and those that are not publicly available as research, for example, drop-out prediction algorithms developed for profit by educational software providers. We plan to work with stakeholders in education, including students, counselors, faculty, and administration, to better understand how the algorithms in our review, and those developed proprietarily, are being implemented in higher education institutions. To move towards a human-centered approach to building evidence-based and theoretically-driven algorithms that do as they intend in real-world scenarios, we plan to seek a better understanding of the connections between input features and proposed interventions, and LMS data and student engagement.

6 CONCLUSION

We conducted a systematic review of existing literature published in the ACM Digital Library on algorithmic decision-making in higher education. After establishing a corpus of 63 papers, we quantitatively analyzed the papers for their input data, computational methods, and target outcomes, before applying an established human-centered algorithm design framework [17]. Going forward, we recommend that the HCI research community focus on developing algorithm design processes with theoretical, participatory, and speculative dimensions. Theoretically-grounded algorithm design that includes active engagement of stakeholders through the design and evaluation phases will ensure algorithms that are better aligned to the socially-complex systems of which they are a part. It will also increase understanding of highly-contextual input data and allow for better interpretation of results. Additionally, we recommend that future research in this area considers the future of higher education institutions and how the proposed algorithms may impact stakeholders.

ACKNOWLEDGMENTS

This research was supported by the Natural Sciences and Engineering Research Council of Canada's Early Career Discovery Grant. We would like to thank the anonymous reviewers for helping us improve this work.

REFERENCES

- [1] Nasir Abdul Jalil and Mikkay Wong Ei Leen. 2021. Learning Analytics in Higher Education: The Student Expectations of Learning Analytics. In 2021 5th International Conference on Education and E-Learning (ICEEL 2021). Association for Computing Machinery, New York, NY, USA, 249–254. https://doi.org/10. 1145/3502434.3502463
- [2] Rediet Abebe, Solon Barocas, Jon Kleinberg, Karen Levy, Manish Raghavan, and David G. Robinson. 2020. Roles for computing in social change. In Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. ACM, Barcelona Spain, 252–260. https://doi.org/10.1145/3351095.3372871
- [3] Everaldo Aguiar, Nitesh V. Chawla, Jay Brockman, G. Alex Ambrose, and Victoria Goodrich. 2014. Engagement vs performance: using electronic portfolios to predict first semester engineering student retention. In Proceedings of the Fourth International Conference on Learning Analytics And Knowledge. ACM, Indianapolis Indiana USA, 103–112. https://doi.org/10.1145/2567574.2567583
- [4] Ritesh Ajoodha, Ashwini Jadhav, and Shalini Dukhan. 2020. Forecasting Learner Attrition for Student Success at a South African University. In Conference of the South African Institute of Computer Scientists and Information Technologists 2020. ACM, Cape Town South Africa, 19–28. https://doi.org/10.1145/3410886.3410973
- [5] Balqis Albreiki, Nazar Zaki, and Hany Alashwal. 2021. A Systematic Literature Review of Student' Performance Prediction Using Machine Learning Techniques. *Education Sciences* 11, 9 (Sept. 2021), 552. https://doi.org/10.3390/ educsci11090552 Number: 9 Publisher: Multidisciplinary Digital Publishing Institute.
- [6] Alejandro Jaimes, Daniel Gatica-Perez, Nicu Sebe, and Thomas S Huang. 2007. Human-Centered Computing: Toward a Human Revolution. Computer (Long Beach, Calif.) 40, 5 (2007), 30–34. https://search.proquest.com/docview/ 197419098?pq-origsite=primo Place: New York Publisher: The Institute of Electrical and Electronics Engineers, Inc. IEEE.
- [7] Amal AlGhamdi, Amal Barsheed, Hanadi AlMshjary, and Hanan AlGhamdi. 2020. A Machine Learning Approach for Graduate Admission Prediction. In Proceedings of the 2020 2nd International Conference on Image, Video and Signal Processing. ACM, Singapore Singapore, 155–158. https://doi.org/10.1145/3388818.3393716
- [8] Saud Altaf, Waseem Soomro, and Mohd Izani Mohamed Rawi. 2019. Student Performance Prediction using Multi-Layers Artificial Neural Networks: A Case Study on Educational Data Mining. In Proceedings of the 2019 3rd International Conference on Information System and Data Mining - ICISDM 2019. ACM Press, Houston, TX, USA, 59–64. https://doi.org/10.1145/3325917.3325919
- [9] Eyman Alyahyan and Dilek Düştegör. 2020. Predicting academic success in higher education: literature review and best practices. *International Journal* of Educational Technology in Higher Education 17, 1 (Dec. 2020), 3. https: //doi.org/10.1186/s41239-020-0177-7

- [10] Renji George Amballoor and Shankar B Naik. 2021. Technological Achievement Index (TAI) in Higher Education- An Empirical Analysis. In 2021 5th International Conference on Digital Technology in Education. ACM, Busan Republic of Korea, 93–96. https://doi.org/10.1145/3488466.3488476
- [11] Sattar Ameri, Mahtab J. Fard, Ratna B. Chinnam, and Chandan K. Reddy. 2016. Survival Analysis based Framework for Early Prediction of Student Dropouts. In Proceedings of the 25th ACM International on Conference on Information and Knowledge Management. ACM, Indianapolis Indiana USA, 903–912. https: //doi.org/10.1145/2983323.2983351
- [12] Cecilia Rodriguez Aragon. 2022. Human-centered data science: an introduction. The MIT Press, Cambridge, Massachusetts.
- [13] Swati Bajpai and S. Mani. 2017. Big Data in Education and Learning Analytics. TechnoLearn: An International Journal of Educational Technology 7, 1and2 (2017), 45. https://doi.org/10.5958/2249-5223.2017.00005.5
- [14] Máté Baranyi, Marcell Nagy, and Roland Molontay. 2020. Interpretable Deep Learning for University Dropout Prediction. In Proceedings of the 21st Annual Conference on Information Technology Education. ACM, Virtual Event USA, 13–19. https://doi.org/10.1145/3368308.3415382
- [15] Rebecca Barber and Mike Sharkey. 2012. Course correction: using analytics to predict course success. In Proceedings of the 2nd International Conference on Learning Analytics and Knowledge. ACM, Vancouver British Columbia Canada, 259–262. https://doi.org/10.1145/2330601.2330664
- [16] Solon Barocas. 2014. Putting Data to Work. In Data and Discrimination: Collected Essays, Seeta Peña Gangadharan, Virginia Eubanks, and Solon Barocas (Eds.). Open Technology Institute. https://timlibert.me/pdf/2014-Data_ Discrimination_Collected_Essays.pdf
- [17] Eric PS Baumer. 2017. Toward human-centered algorithm design. Big Data & Society 4, 2 (Dec. 2017), 2053951717718854. https://doi.org/10.1177/ 2053951717718854 Publisher: SAGE Publications Ltd.
- [18] Ceasar Ian P. Benablo, Evangeline T. Sarte, Joe Marie D. Dormido, and Thelma Palaoag. 2018. Higher Education Student's Academic Performance Analysis through Predictive Analytics. In Proceedings of the 2018 7th International Conference on Software and Computer Applications. ACM, Kuantan Malaysia, 238–242. https://doi.org/10.1145/3185089.3185102
- [19] Inma Borrella, Sergio Caballero-Caballero, and Eva Ponce-Cueto. 2019. Predict and Intervene: Addressing the Dropout Problem in a MOOC-based Program. In Proceedings of the Sixth (2019) ACM Conference on Learning @ Scale. ACM, Chicago IL USA, 1–9. https://doi.org/10.1145/3330430.3333634
- [20] Nynke Bos and Saskia Brand-Gruwel. 2016. Student differences in regulation strategies and their use of learning resources: implications for educational design. In Proceedings of the Sixth International Conference on Learning Analytics & Knowledge - LAK '16. ACM Press, Edinburgh, United Kingdom, 344–353. https://doi.org/10.1145/2883851.2883890
- [21] Christopher Brooks and Jim Greer. 2014. Explaining predictive models to learning specialists using personas. In Proceedings of the Fourth International Conference on Learning Analytics And Knowledge (LAK '14). Association for Computing Machinery, New York, NY, USA, 26–30. https://doi.org/10.1145/ 2567574.2567512
- [22] John P. Buerck, Srikanth P. Mudigonda, Stephanie E. Mooshegian, Kyle Collins, Nicholas Grimm, Kristen Bonney, and Hadley Kombrink. 2013. Predicting nontraditional student learning outcomes using data analytics - a pilot research study. *Journal of Computing Sciences in Colleges* 28, 5 (May 2013), 260–265.
- [23] Joey A. Cabrera, Markdy Y. Orong, Nelpa N. Capio, Arnel Filarca, Eden Neri, and Ariel R. Clarin. 2020. A Data Mining Approach for Student Referral Service of the Guidance Center: An Input in Designing Mediation Scheme for Higher Education Institutions of the Philippines. In Proceedings of the 3rd International Conference on Software Engineering and Information Management. ACM, Sydney NSW Australia, 10–14. https://doi.org/10.1145/3378936.3378958
- [24] Arthur Ree Campbell and Charlie J. Dickson. 1996. Predicting Student Success: A 10-Year Review Using Integrative Review and Meta-Analysis. *Journal of Professional Nursing* 12, 1 (1996), 47–59.
- [25] Stevie Chancellor, Eric P. S. Baumer, and Munmun De Choudhury. 2019. Who is the "Human" in Human-Centered Machine Learning: The Case of Predicting Mental Health from Social Media. Proceedings of the ACM on Human-Computer Interaction 3, CSCW (Nov. 2019), 1–32. https://doi.org/10.1145/3359249
- [26] Youjie Chen, Annie Fu, Jennifer Jia-Ling Lee, Ian Wilkie Tomasik, and René F. Kizilcec. 2022. Pathways: Exploring Academic Interests with Historical Course Enrollment Records. In Proceedings of the Ninth ACM Conference on Learning @ Scale. ACM, New York City NY USA, 222–233. https://doi.org/10.1145/3491140. 3528270
- [27] Marianne Cherrington, David Airehrour, Joan Lu, Qiang Xu, David Cameron-Brown, and Ihaka Dunn. 2020. Features of Human-Centred Algorithm Design. In 2020 30th International Telecommunication Networks and Applications Conference (ITNAC). 1–6. https://doi.org/10.1109/ITNAC50341.2020.9315169 ISSN: 2474-154X.
- [28] Shruthi Chockkalingam, Run Yu, and Zachary A. Pardos. 2021. Which one's more work? Predicting effective credit hours between courses. In LAK21: 11th International Learning Analytics and Knowledge Conference. ACM, Irvine CA

USA, 599-605. https://doi.org/10.1145/3448139.3448204

- [29] Sylvia Chong, Yew Haur Lee, and Yoke Wah Tang. 2020. Data Analytics and Visualization to Support the Adult Learner in Higher Education. In 2020 The 4th International Conference on E-Society, E-Education and E-Technology. ACM, Taipei Taiwan, 126–131. https://doi.org/10.1145/3421682.3421698
- [30] Ali Daud, Naif Radi Aljohani, Rabeeh Ayaz Abbasi, Miltiadis D. Lytras, Farhat Abbas, and Jalal S. Alowibdi. 2017. Predicting Student Performance using Advanced Learning Analytics. In Proceedings of the 26th International Conference on World Wide Web Companion - WWW '17 Companion. ACM Press, Perth, Australia, 415–421. https://doi.org/10.1145/3041021.3054164
- [31] Shane Dawson, Jelena Jovanovic, Dragan Gašević, and Abelardo Pardo. 2017. From prediction to impact: evaluation of a learning analytics retention program. In Proceedings of the Seventh International Learning Analytics & Knowledge Conference. ACM, Vancouver British Columbia Canada, 474–478. https://doi. org/10.1145/3027385.3027405
- [32] Cristobal de Brey, Lauren Musu, Joel McFarland, Sidney Wilkinson-Flicker, Melissa Diliberti, Anlan Zhang, Claire Branstetter, and Xiaolei Wang. 2019. Status and Trends in the Education of Racial and Ethnic Groups 2018. Technical Report. U.S. Department of Education, National Center for Education Statistics, Institute of Education Sciences. 228 pages.
- [33] Fernando Delgado, Solon Barocas, and Karen Levy. 2022. An Uncommon Task: Participatory Design in Legal AI. Proceedings of the ACM on Human-Computer Interaction 6, CSCW1 (April 2022), 51:1–51:23. https://doi.org/10.1145/3512898
- [34] Carrie Demmans Epp, Krystle Phirangee, Jim Hewitt, and Charles A. Perfetti. 2020. Learning management system and course influences on student actions and learning experiences. *Educational Technology Research and Development* 68, 6 (Dec. 2020), 3263–3297. https://doi.org/10.1007/s11423-020-09821-1
- [35] Vanessa Echeverría, Allan Avendaño, Katherine Chiluiza, Aníbal Vásquez, and Xavier Ochoa. 2014. Presentation Skills Estimation Based on Video and Kinect Data Analysis. In Proceedings of the 2014 ACM workshop on Multimodal Learning Analytics Workshop and Grand Challenge. ACM, Istanbul Turkey, 53–60. https: //doi.org/10.1145/2666633.2666641
- [36] Rebecca L. Edwards, Sarah K. Davis, Allyson F. Hadwin, and Todd M. Milford. 2017. Using predictive analytics in a self-regulated learning university course to promote student success. In Proceedings of the Seventh International Learning Analytics & Knowledge Conference. ACM, Vancouver British Columbia Canada, 556–557. https://doi.org/10.1145/3027385.3029455
- [37] Hamza Errahmouni Barkam, Max Wang, Barbara Martinez Neda, and Sergio Gago-Masague. 2022. Testing Machine Learning Models to Identify Computer Science Students at High-risk of Probation. In Proceedings of the 53rd ACM Technical Symposium on Computer Science Education V. 2. ACM, Providence RI USA, 1161–1161. https://doi.org/10.1145/3478432.3499103
- [38] Jacqueline Feild, Nicholas Lewkow, Sean Burns, and Karen Gebhardt. 2018. A generalized classifier to identify online learning tool disengagement at scale. In Proceedings of the 8th International Conference on Learning Analytics and Knowledge. ACM, Sydney New South Wales Australia, 61–70. https://doi.org/ 10.1145/3170358.3170370
- [39] Stefano Fiorini, Adrienne Sewell, Mathew Bumbalough, Pallavi Chauhan, Linda Shepard, George Rehrey, and Dennis Groth. 2018. An application of participatory action research in advising-focused learning analytics. In Proceedings of the 8th International Conference on Learning Analytics and Knowledge. ACM, Sydney New South Wales Australia, 89–96. https://doi.org/10.1145/3170358.3170387
- [40] Asbjørn Ammitzbøll Flügge. 2020. Algorithmic Decision Making in Public Administration: A CSCW-Perspective. In Companion of the 2020 ACM International Conference on Supporting Group Work (GROUP '20). Association for Computing Machinery, New York, NY, USA, 15–24. https://doi.org/10.1145/3323994.3371016
- [41] Sarah E. Fox, Vera Khovanskaya, Clara Crivellaro, Niloufar Salehi, Lynn Dombrowski, Chinmay Kulkarni, Lilly Irani, and Jodi Forlizzi. 2020. Worker-Centered Design: Expanding HCI Methods for Supporting Labor. In Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems (CHI EA '20). Association for Computing Machinery, New York, NY, USA, 1–8. https://doi.org/10.1145/3334480.3375157
- [42] Eslam Abou Gamie, M. Samir Abou El-Seoud, and Mostafa A. Salama. 2019. A layered-analysis of the features in higher education data set. In Proceedings of the 2019 8th International Conference on Software and Information Engineering. ACM, Cairo Egypt, 237–242. https://doi.org/10.1145/3328833.3328850
- [43] Twinkle Mae C. Gatbonton and Betchie E. Aguinaldo. 2018. Employability Predictive Model Evaluator Using PART and JRip Classifier. In Proceedings of the 6th International Conference on Information Technology: IoT and Smart City -ICIT 2018. ACM Press, Hong Kong, Hong Kong, 307–310. https://doi.org/10. 1145/3301551.3301569
- [44] Isabelle Guyon and André Elisseeff. 2003. An introduction to variable and feature selection. *The Journal of Machine Learning Research* 3, null (March 2003), 1157–1182.
- [45] Patrick Hall. 2012. Incident 135: University of Texas at Austin's Algorithm to Evaluate Graduate Applications, GRADE, Allegedly Exacerbated Existing Inequality for Marginalized Applicants, Prompting Tool Suspension. https: //incidentdatabase.ai/cite/135

- [46] Patrick Hall. 2020. Incident 140: ProctorU's Identity Verification and Exam Monitoring Systems Provided Allegedly Discriminatory Experiences for BIPOC Students. https://incidentdatabase.ai/cite/140
- [47] Judith M. Harackiewicz and Stacy J. Priniski. 2018. Improving Student Outcomes in Higher Education: The Science of Targeted Intervention. *Annual review of psychology* 69 (Jan. 2018), 409–435. https://doi.org/10.1146/annurev-psych-122216-011725
- [48] Arto Hellas, Petri Ihantola, Andrew Petersen, Vangel V. Ajanovski, Mirela Gutica, Timo Hynninen, Antti Knutas, Juho Leinonen, Chris Messom, and Soohyun Nam Liao. 2018. Predicting academic performance: a systematic literature review. In Proceedings Companion of the 23rd Annual ACM Conference on Innovation and Technology in Computer Science Education. ACM, Larnaca Cyprus, 175–199. https://doi.org/10.1145/3293881.3295783
- [49] Ho Thao Hien, Pham-Nguyen Cuong, Le Nguyen Hoai Nam, Ho Le Thi Kim Nhung, and Le Dinh Thang. 2018. Intelligent Assistants in Higher-Education Environments: The FIT-EBot, a Chatbot for Administrative and Learning Support. In Proceedings of the Ninth International Symposium on Information and Communication Technology - SoICT 2018. ACM Press, Danang City, Viet Nam, 69–76. https://doi.org/10.1145/3287921.3287937
- [50] Martin Hlosta, Zdenek Zdrahal, and Jaroslav Zendulka. 2017. Ouroboros: early identification of at-risk students without models based on legacy data. In Proceedings of the Seventh International Learning Analytics & Knowledge Conference. ACM, Vancouver British Columbia Canada, 6–15. https://doi.org/10.1145/ 3027385.3027449
- [51] Sandeep M. Jayaprakash and Eitel J. M. Lauría. 2014. Open academic early alert system: technical demonstration. In Proceedings of the Fourth International Conference on Learning Analytics And Knowledge. ACM, Indianapolis Indiana USA, 267–268. https://doi.org/10.1145/2567574.2567578
- [52] Weijie Jiang and Zachary A. Pardos. 2019. Time slice imputation for personalized goal-based recommendation in higher education. In *Proceedings of the 13th ACM Conference on Recommender Systems*. ACM, Copenhagen Denmark, 506–510. https://doi.org/10.1145/3298689.3347030
- [53] Weijie Jiang and Zachary A. Pardos. 2021. Towards Equity and Algorithmic Fairness in Student Grade Prediction. In Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society. ACM, Virtual Event USA, 608-617. https: //doi.org/10.1145/3461702.3462623
- [54] George Karypis. 2017. Improving Higher Education: Learning Analytics & Recommender Systems Research. In Proceedings of the Eleventh ACM Conference on Recommender Systems (RecSys '17). Association for Computing Machinery, New York, NY, USA, 2. https://doi.org/10.1145/3109859.3109870
- [55] Prableen Kaur, Agoritsa Polyzou, and George Karypis. 2019. Causal Inference in Higher Education: Building Better Curriculums. In Proceedings of the Sixth (2019) ACM Conference on Learning @ Scale. ACM, Chicago IL USA, 1–4. https: //doi.org/10.1145/3330430.3333663
- [56] Seunghyun Kim, Afsaneh Razi, Gianluca Stringhini, Pamela J. Wisniewski, and Munmun De Choudhury. 2021. A Human-Centered Systematic Literature Review of Cyberbullying Detection Algorithms. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (Oct. 2021), 1–34. https://doi.org/10. 1145/3476066
- [57] Georgios Kostopoulos, Sotiris Kotsiantis, and Panagiotis Pintelas. 2015. Estimating student dropout in distance higher education using semi-supervised techniques. In Proceedings of the 19th Panhellenic Conference on Informatics. ACM, Athens Greece, 38–43. https://doi.org/10.1145/2801948.2802013
- [58] Mukesh Kumar, A.J. Singh, and Disha Handa. 2017. Literature Survey on Student's Performance Prediction in Education using Data Mining Techniques. *International Journal of Education and Management Engineering* 7, 6 (Nov. 2017), 40–49. https://doi.org/10.5815/ijeme.2017.06.05
- [59] Catherine Kung and Renzhe Yu. 2020. Interpretable Models Do Not Compromise Accuracy or Fairness in Predicting College Success. In Proceedings of the Seventh ACM Conference on Learning @ Scale. ACM, Virtual Event USA, 413–416. https: //doi.org/10.1145/3386527.3406755
- [60] Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual Fairness. In Advances in Neural Information Processing Systems, Vol. 30. Curran Associates, Inc. https://proceedings.neurips.cc/paper/2017/ hash/a486cd07e4ac3d270571622f4f316ec5-Abstract.html
- [61] David Lang, Alex Wang, Nathan Dalal, Andreas Paepcke, and Mitchell Stevens. 2021. Forecasting Undergraduate Majors Using Academic Transcript Data. In Proceedings of the Eighth ACM Conference on Learning @ Scale. ACM, Virtual Event Germany, 243–246. https://doi.org/10.1145/3430895.3460149
- [62] Hansol Lee and René F. Kizilcec. 2020. Evaluation of Fairness Trade-offs in Predicting Student Success. Technical Report arXiv:2007.00088. arXiv. https: //doi.org/10.48550/arXiv.2007.00088 arXiv:2007.00088 [cs] type: article.
- [63] Chenglu Li, Wanli Xing, and Walter Leite. 2021. Yet Another Predictive Model? Fair Predictions of Students' Learning Outcomes in an Online Math Learning Platform. In LAK21: 11th International Learning Analytics and Knowledge Conference. ACM, Irvine CA USA, 572–578. https://doi.org/10.1145/3448139.3448200
- [64] Gonzalo Luzardo, Bruno Guamán, Katherine Chiluiza, Jaime Castells, and Xavier Ochoa. 2014. Estimation of Presentations Skills Based on Slides and Audio

Features. In Proceedings of the 2014 ACM workshop on Multimodal Learning Analytics Workshop and Grand Challenge. ACM, Istanbul Turkey, 37–44. https: //doi.org/10.1145/2666633.2666639

- [65] Xiaofeng Ma, Yan Yang, and Zhurong Zhou. 2018. Using Machine Learning Algorithm to Predict Student Pass Rates In Online Education. In Proceedings of the 3rd International Conference on Multimedia Systems and Signal Processing -ICMSSP '18. ACM Press, Shenzhen, China, 156–161. https://doi.org/10.1145/ 3220162.3220188
- [66] Charlot L. Maramag and Thelma D. Palaoag. 2019. Assessing CSU Students' Academic Performance on iLearn Portal Using Data Analytics. In Proceedings of the 2019 5th International Conference on Computing and Artificial Intelligence -ICCAI '19. ACM Press, Bali, Indonesia, 25–29. https://doi.org/10.1145/3330482. 3330495
- [67] Frank Marcinkowski, Kimon Kieslich, Christopher Starke, and Marco Lünich. 2020. Implications of AI (un-)fairness in higher education admissions: the effects of perceived AI (un-)fairness on exit, voice and organizational reputation. In Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20). Association for Computing Machinery, New York, NY, USA, 122–130. https://doi.org/10.1145/3351095.3372867
- [68] Jen McPherson, Huong Ly Tong, Scott J. Fatt, and Danny Y. T. Liu. 2016. Student perspectives on data provision and use: starting to unpack disciplinary differences. In Proceedings of the Sixth International Conference on Learning Analytics & Knowledge (LAK '16). Association for Computing Machinery, New York, NY, USA, 158–167. https://doi.org/10.1145/2883851.2883945
- [69] David Moher, Alessandro Liberati, Jennifer Tetzlaff, and Douglas G. Altman. 2009. Preferred reporting items for systematic reviews and meta-analyses: the PRISMA statement. BMJ 339 (July 2009), b2535. https://doi.org/10.1136/bmj.b2535 Publisher: British Medical Journal Publishing Group Section: Research Methods & amp; Reporting.
- [70] John T. Mouw and Ritu K. Khanna. 1993. Prediction of academic success: A review of the literature and some recommendations. *College Student Journal* 27, 3 (1993), 328–336. Place: US Publisher: Project Innovation of Mobile.
- [71] Arham Muslim, Mohamed Amine Chatti, Tanmaya Mahapatra, and Ulrik Schroeder. 2016. A rule-based indicator definition tool for personalized learning analytics. In Proceedings of the Sixth International Conference on Learning Analytics & Knowledge - LAK '16. ACM Press, Edinburgh, United Kingdom, 264–273. https://doi.org/10.1145/2883851.2883921
- [72] Celia González Nespereira, Kais Dai, Rebeca P. Díaz Redondo, and Ana Fernández Vilas. 2014. Is the LMS access frequency a sign of students' success in face-to-face higher education?. In Proceedings of the Second International Conference on Technological Ecosystems for Enhancing Multiculturality - TEEM '14. ACM Press, Salamanca, Spain, 283–290. https://doi.org/10.1145/2669711.2669912
- [73] Charbel Obeid, Inaya Lahoud, Hicham El Khoury, and Pierre-Antoine Champin. 2018. Ontology-based Recommender System in Higher Education. In Companion of the The Web Conference 2018 on The Web Conference 2018 - WWW '18. ACM Press, Lyon, France, 1031–1034. https://doi.org/10.1145/3184558.3191533
- [74] Francis Ofori, Elizaphan Maina, and Rhoda Gitonga. 2020. Using Machine Learning Algorithms to Predict Studentsâ€[™] Performance and Improve Learning Outcome: A Literature Based Review. *Journal of Information and Technology* 4, 1 (March 2020). https://stratfordjournals.org/journals/index.php/Journal-of-Information-and-Techn/article/view/480 Number: 1.
- [75] Opeyemi Ojajuni, Foluso Ayeni, Olagunju Akodu, Femi Ekanoye, Samson Adewole, Timothy Ayo, Sanjay Misra, and Victor Mbarika. 2021. Predicting Student Academic Performance Using Machine Learning. In Computational Science and Its Applications ICCSA 2021 (Lecture Notes in Computer Science), Osvaldo Gervasi, Beniamino Murgante, Sanjay Misra, Chiara Garau, Ivan Blečić, David Taniar, Bernady O. Apduhan, Ana Maria A. C. Rocha, Eufemia Tarantino, and Carmelo Maria Torre (Eds.). Springer International Publishing, Cham, 481–491. https://doi.org/10.1007/978-3-030-87013-3_36
- [76] Mark Olssen * and Michael A. Peters. 2005. Neoliberalism, higher education and the knowledge economy: from the free market to knowledge capitalism. *Journal of Education Policy* 20, 3 (Jan. 2005), 313–345. https://doi.org/10.1080/ 02680930500108718
- [77] Svyatoslav Oreshin, Andrey Filchenkov, Polina Petrusha, Egor Krasheninnikov, Alexander Panfilov, Igor Glukhov, Yulia Kaliberda, Daniil Masalskiy, Alexey Serdyukov, Vladimir Kazakovtsev, Maksim Khlopotov, Timofey Podolenchuk, Ivan Smetannikov, and Daria Kozlova. 2020. Implementing a Machine Learning Approach to Predicting Students' Academic Outcomes. In 2020 International Conference on Control, Robotics and Intelligent System. ACM, Xiamen China, 78–83. https://doi.org/10.1145/3437802.3437816
- [78] Markdy Y. Orong, Roseclaremath A. Caroro, Geraldine D. Durias, Joey A. Cabrera, Herwina Lonzon, and Gretel T. Ricalde. 2020. A Predictive Analytics Approach in Determining the Predictors of Student Attrition in the Higher Education Institutions in the Philippines. In *Proceedings of the 3rd International Conference on Software Engineering and Information Management*. ACM, Sydney NSW Australia, 222–225. https://doi.org/10.1145/3378936.3378956
- [79] David Otoo-Arthur and Terence Van Zyl. 2019. A Systematic Review on Big Data Analytics Frameworks for Higher Education - Tools and Algorithms. In

Proceedings of the 2019 2nd International Conference on E-Business, Information Management and Computer Science. ACM, Kuala Lumpur Malaysia, 1–9. https: //doi.org/10.1145/3377817.3377836

- [80] Zachary A. Pardos, Hung Chau, and Haocheng Zhao. 2019. Data-Assistive Course-to-Course Articulation Using Machine Translation. In Proceedings of the Sixth (2019) ACM Conference on Learning @ Scale. ACM, Chicago IL USA, 1–10. https://doi.org/10.1145/3330430.3333622
- [81] Zachary A. Pardos and Weijie Jiang. 2020. Designing for serendipity in a university course recommendation system. In *Proceedings of the Tenth International Conference on Learning Analytics & Knowledge*. ACM, Frankfurt Germany, 350–359. https://doi.org/10.1145/3375462.3375524
- [82] Mar Perez-Sanagustin, Ronald Pérez-Álvarez, Jorge Maldonado-Mahauad, Esteban Villalobos, Isabel Hilliger, Josefina Hernández, Diego Sapunar, Pedro Manuel Moreno-Marcos, Pedro J. Muñoz-Merino, Carlos Delgado Kloos, and Jon Imaz. 2021. Can Feedback based on Predictive Data Improve Learners' Passing Rates in MOOCs? A Preliminary Analysis. In Proceedings of the Eighth ACM Conference on Learning @ Scale. ACM, Virtual Event Germany, 339–342. https: //doi.org/10.1145/3430895.3460991
- [83] Kian L. Pokorny. 2021. A machine learning approach to understanding the viability of private 4-year higher-education institutions. *Journal of Computing Sciences in Colleges* 37, 4 (Oct. 2021), 50–57.
- [84] Paul Prinsloo and Sharon Slade. 2017. An elephant in the learning analytics room: the obligation to act. In Proceedings of the Seventh International Learning Analytics & Knowledge Conference (LAK '17). Association for Computing Machinery, New York, NY, USA, 46–55. https://doi.org/10.1145/3027385.3027406
- [85] Afsaneh Razi, Seunghyun Kim, Ashwaq Alsoubai, Gianluca Stringhini, Thamar Solorio, Munmun De Choudhury, and Pamela J. Wisniewski. 2021. A Human-Centered Systematic Literature Review of the Computational Approaches for Online Sexual Risk Detection. Proceedings of the ACM on Human-Computer Interaction 5, CSCW2 (Oct. 2021), 465:1-465:38. https://doi.org/10.1145/3479609
- [86] Nova Rijati, Surya Sumpeno, and Mauridhi Hery Purnomo. 2018. Multi-Attribute Clustering of Student's Entrepreneurial Potential Mapping Based on Its Characteristics and the Affecting Factors: Preliminary Study on Indonesian Higher Education Database. In Proceedings of the 2018 10th International Conference on Computer and Automation Engineering. ACM, Brisbane Australia, 11–16. https://doi.org/10.1145/3192975.3193014
- [87] Tim Rogers, Cassandra Colvin, and Belinda Chiera. 2014. Modest analytics: using the index method to identify students at risk of failure. In Proceedings of the Fourth International Conference on Learning Analytics And Knowledge. ACM, Indianapolis Indiana USA, 118–122. https://doi.org/10.1145/2567574.2567629
- [88] Cynthia Rudin. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* 1, 5 (May 2019), 206–215. https://doi.org/10.1038/s42256-019-0048-x
- [89] Lucas H. Sallaberry, Romero Tori, and Fatima L S Nunes. 2021. Comparison of machine learning algorithms for automatic assessment of performance in a virtual reality dental simulator. In Symposium on Virtual and Augmented Reality. ACM, Virtual Event Brazil, 14–23. https://doi.org/10.1145/3488162.3488207
- [90] Lidia Sandra, Ford Lumbangaol, and Tokuro Matsuo. 2021. Machine Learning Algorithm to Predict Student's Performance: A Systematic Literature Review. *TEM Journal* (Nov. 2021), 1919–1927. https://doi.org/10.18421/TEM104-56
- [91] Devansh Saxena, Karla Badillo-Urquiola, Pamela J. Wisniewski, and Shion Guha. 2020. A Human-Centered Review of Algorithms used within the U.S. Child Welfare System. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems. ACM, Honolulu HI USA, 1–15. https://doi.org/10.1145/ 3313831.3376229
- [92] Devansh Saxena and Shion Guha. 2020. Conducting Participatory Design to Improve Algorithms in Public Services: Lessons and Challenges. In Conference Companion Publication of the 2020 on Computer Supported Cooperative Work and Social Computing (CSCW '20 Companion). Association for Computing Machinery, New York, NY, USA, 383–388. https://doi.org/10.1145/3406865.3418331
- [93] Boran Sekeroglu, Kamil Dimililer, and Kubra Tuncal. 2019. Student Performance Prediction and Classification Using Machine Learning Algorithms. In Proceedings of the 2019 8th International Conference on Educational and Information Technology. ACM, Cambridge United Kingdom, 7–11. https://doi.org/10.1145/ 3318396.3318419
- [94] Richard J. Self. 2014. Governance Strategies for the Cloud, Big Data, and Other Technologies in Education. In Proceedings of the 2014 IEEE/ACM 7th International Conference on Utility and Cloud Computing (UCC '14). IEEE Computer Society, USA, 630–635. https://doi.org/10.1109/UCC.2014.101
- [95] Amirah Mohamed Shahiri, Wahidah Husain, and Nur'aini Abdul Rashid. 2015. A Review on Predicting Student's Performance Using Data Mining Techniques. *Procedia Computer Science* 72 (Jan. 2015), 414–422. https://doi.org/10.1016/j. procs.2015.12.157
- [96] Ben Shneiderman. 2021. Human-Centered AI. Issues in Science and Technology 37, 2 (2021), 56–61. https://www.proquest.com/docview/2481224128? parentSessionId=chnCbaQ2NZxuTZjdR%2FSmni%2BibODmNifp1bKKqe% 2FEjkA%3D&pq-origsite=primo& Num Pages: 56-61 Place: Washington, United

States Publisher: Issues in Science and Technology.

- [97] Ben Shneiderman. 2022. Human-Centered AI. Oxford University Press, Oxford, New York.
- [98] Sharon Slade and Paul Prinsloo. 2013. Learning Analytics: Ethical Issues and Dilemmas. American Behavioral Scientist 57, 10 (Oct. 2013), 1510–1529. https: //doi.org/10.1177/0002764213479366 Publisher: SAGE Publications Inc.
- [99] Farahnaz Soleimani and Jeonghyun Lee. 2021. Comparative Analysis of the Feature Extraction Approaches for Predicting Learners Progress in Online Courses: MicroMasters Credential versus Traditional MOOCs. In Proceedings of the Eighth ACM Conference on Learning @ Scale (L@S '21). Association for Computing Machinery, New York, NY, USA, 151–159. https://doi.org/10.1145/3430895.3460143
- [100] Shawn Staudaher, Jeonghyun Lee, and Farahnaz Soleimani. 2020. Predicting Applicant Admission Status for Georgia Tech's Online Master's in Analytics Program. In Proceedings of the Seventh ACM Conference on Learning @ Scale. ACM, Virtual Event USA, 309–312. https://doi.org/10.1145/3386527.3406735
- [101] Nick Stockton. 2020. Incident 78: Meet the Secret Algorithm That's Keeping Students Out of College. https://incidentdatabase.ai/cite/78
- [102] Evis Trandafili, Alban Allkoçi, Elinda Kajo, and Aleksandër Xhuvani. 2012. Discovery and evaluation of student's profiles with machine learning. In Proceedings of the Fifth Balkan Conference in Informatics on - BCI '12. ACM Press, Novi Sad, Serbia, 174. https://doi.org/10.1145/2371316.2371350
- [103] Steven Van Goidsenhoven, Daria Bogdanova, Galina Deeva, Seppe vanden Broucke, Jochen De Weerdt, and Monique Snoeck. 2020. Predicting student success in a blended learning environment. In Proceedings of the Tenth International Conference on Learning Analytics & Knowledge. ACM, Frankfurt Germany, 17-25. https://doi.org/10.1145/3375462.3375494
- [104] Jennifer Wortman Vaughan and Hanna Wallach. 2022. A Human-Centered Agenda for Intelligible Machine Learning. (Sept. 2022). https://www.microsoft.com/en-us/research/publication/a-human-centeredagenda-for-intelligible-machine-learning/
- [105] Weichen Wang, Subigya Nepal, Jeremy F. Huckins, Lessley Hernandez, Vlado Vojdanovski, Dante Mack, Jane Plomp, Arvind Pillai, Mikio Obuchi, Alex daSilva, Eilis Murphy, Elin Hedlund, Courtney Rogers, Meghan Meyer, and Andrew Campbell. 2022. First-Gen Lens: Assessing Mental Health of First-Generation Students across Their First Year at College Using Mobile Sensing. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies 6, 2 (July 2022), 1–32. https://doi.org/10.1145/3543194
- [106] Drew Wham. 2017. Forecasting student outcomes at university-wide scale using machine learning. In Proceedings of the Seventh International Learning Analytics & Knowledge Conference. ACM, Vancouver British Columbia Canada, 576–577. https://doi.org/10.1145/3027385.3029467
- [107] Kimberly Williamson and René F. Kizilcec. 2021. Learning Analytics Dashboard Research Has Neglected Diversity, Equity and Inclusion. In Proceedings of the Eighth ACM Conference on Learning @ Scale (L@S '21). Association for Computing Machinery, New York, NY, USA, 287–290. https://doi.org/10.1145/3430895. 3460160
- [108] W. K. Winters. 1971. A scheduling algorithm for a computer assisted registration system. Commun. ACM 14, 3 (March 1971), 166–171. https://doi.org/10.1145/ 362566.362569
- [109] Jacob O. Wobbrock and Julie A. Kientz. 2016. Research contributions in humancomputer interaction. *Interactions* 23, 3 (April 2016), 38–44. https://doi.org/10. 1145/2907069
- [110] Hongli Yan. 2021. The Trends and Challenges of Emerging Technologies in Higher Education. In 2021 2nd International Conference on Education Development and Studies. ACM, Hilo HI USA, 89–95. https://doi.org/10.1145/3459043.3459060
- [111] Renzhe Yu, Hansol Lee, and René F. Kizilcec. 2021. Should College Dropout Prediction Models Include Protected Attributes?. In Proceedings of the Eighth ACM Conference on Learning @ Scale. ACM, Virtual Event Germany, 91–100. https://doi.org/10.1145/3430895.3460139
- [112] Renzhe Yu, Qiujie Li, and Christian Fischer. 2020. Towards Accurate and Fair Prediction of College Success: Evaluating Different Sources of Student Data. In Educational Data Mining. Online, 10.
- [113] Taeho Yu and Il-Hyun Jo. 2014. Educational technology approach toward learning analytics: relationship between student online behavior and learning performance in higher education. In Proceedings of the Fourth International Conference on Learning Analytics And Knowledge. ACM, Indianapolis Indiana USA, 269–270. https://doi.org/10.1145/2567574.2567594
- [114] Nastaran Zanjani. 2017. The important elements of LMS design that affect user engagement with e-learning tools within LMSs in the higher education sector. Australasian Journal of Educational Technology 33, 1 (April 2017). https: //doi.org/10.14742/ajet.2938 Number: 1.
- [115] 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 (Dec. 2019), 39. https://doi.org/10.1186/s41239-019-0171-0
- [116] Hassan Zeineddine, Udo Braendle, and Assaad Farah. 2019. Auto-generated ensemble model for predicting student success. In Proceedings of the Second

International Conference on Data Science, E-Learning and Information Systems -DATA '19. ACM Press, Dubai, United Arab Emirates, 1–4. https://doi.org/10. 1145/3368691.3368714

- [117] Yupei Zhang, Rui An, Jiaqi Cui, and Xuequn Shang. 2021. Undergraduate Grade Prediction in Chinese Higher Education Using Convolutional Neural Networks. In LAK21: 11th International Learning Analytics and Knowledge Conference. ACM, Irvine CA USA, 462–468. https://doi.org/10.1145/3448139.3448184
- [118] Qing Zhou, Qiang Zhang, and Hao Li. 2021. Analysis on the level of higher universities in different countries using entropy weight method and analytic hierarchy process. In 2021 4th International Conference on Information Systems and Computer Aided Education. ACM, Dalian China, 2722–2727. https://doi. org/10.1145/3482632.3487502

A CODE SHEET

Reference	Input Data	Target Output	Computational Methods	Study Design
[3]	Demographic; Grade/GPA; Protected Class	Retention	ML	
[4]	Demographic; Grade/GPA; Protected Class	Retention	ML	
[7]	Grade/GPA	Admissions	ML	
[8]	LMS/Engagement; Grade/GPA	Retention	DL	
[10]	Institutional	Institutional Planning	Statistical Methods	
[11]	Demographic; Grade/GPA; Enrollment/Pathways; Protected Class	Retention	Statistical Methods	
[14]	Demographic; Grade/GPA; Enrollment/Pathways; Protected Class	Retention	ML; DL	
[15]	Demographic; LMS/Engagement; Grade/GPA; Protected Class	Retention	Statistical Methods	
[18]	Demographic; Grade/GPA; Protected Class	Grade Prediction	ML	
[19]	LMS/Engagement; Grade/GPA; Enrollment/Pathways	Retention	ML	Theoretical Design
[20]	LMS/Engagement; Student Survey	Grade Prediction	Statistical Methods	
[22]	LMS/Engagement; Grade/GPA	Grade Prediction	Statistical Methods	
[23]	Demographic; Protected Class	Student Services	Statistical Methods	
[26]	Enrollment/Pathways	Pathway Advising	NLP	Theoretical Design, Participatory Design
[28]	Grade/GPA; Enrollment/Pathways	Institutional Planning	ML; DL; NLP	
[29]	Demographic; Grade/GPA; Protected Class	Grade Prediction	ML; DL	
[30]	Demographic; Protected Class	Retention	ML	
[31]	Demographic; Protected Class	Retention	Statistical Methods	
[35]		Assessment	ML	
[36]	LMS/Engagement; Grade/GPA; Student Survey	Grade Prediction	Statistical Methods	
[37]	Demographic; Grade/GPA; Protected Class	Retention	ML	
[38]	Grade/GPA	Engagement	ML	
[39]	Grade/GPA; Enrollment/Pathways	Retention	ML;	Participatory Design
[42]	Demographic; LMS/Engagement; Grade/GPA; Enrollment/Pathways; Protected Class	Engagement	ML	
[43]	Grade/GPA	Institutional Planning	Rules-Based	
[49]	Student Survey	Student Services	DL; NLP	
[50]	Demographic; LMS/Engagement; Protected Class	Retention	ML	
[51]	Demographic; LMS/Engagement; Grade/GPA; Protected Class	Retention	ML	
[52]	LMS/Engagement; Grade/GPA; Enrollment/Pathways	Pathway Advising	DL	
[53]	Demographic; Grade/GPA; Enrollment/Pathways; Protected Class	Grade Prediction	DL	Speculative Design
[55]	Grade/GPA	Pathway Advising	Statistical Methods	
[57]	Demographic; Protected Class	Retention	ML	
[59]	Demographic; LMS/Engagement; Grade/GPA; Protected Class	Grade Prediction	ML	Speculative Design
[61]	Enrollment/Pathways	Pathway Advising	ML	
[64]		Assessment	ML	
[65]	Demographic; Grade/GPA; Protected Class	Grade Prediction	ML; DL	
[66]	Demographic; Student Survey; Protected Class	Retention	Statistical Methods	
[71]	LMS/Engagement	Engagement	Rules-Based	
[72]	LMS/Engagement	Grade Prediction	Statistical Methods; ML	
[73]	Student Survey	Pathway Advising	ML	
[75]	Demographic; Grade/GPA; Protected Class	Grade Prediction	ML	
[77]	Demographic; Protected Class	Retention	ML	
[78]	Demographic; Protected Class	Retention	ML	
[80]	Enrollment/Pathways	Pathway Advising	ML; NLP	
[81]	Enrollment/Pathways	Pathway Advising	DL; NLP	Participatory Design
[82]	LMS/Engagement	Retention	ML;	
[83]	Institutional; Enrollment/Pathways	Institutional Planning	ML	
[86]	Demographic; Grade/GPA	Student Services	ML	
[87]	Demographic; LMS/Engagement; Grade/GPA; Enrollment/Pathways; Protected Class	Grade Prediction	Statistical Methods	
[89]	-	Assessment	ML; DL	
[93]	Grade/GPA	Grade Prediction	ML; DL	
[99]	LMS/Engagement	Engagement	ML;	
[100]	Demographic; Protected Class	Admissions	ML	
[102]	Grade/GPA; Enrollment/Pathways	Grade Prediction	ML;	
[103]	LMS/Engagement	Grade Prediction	ML	
[105]	Student Survey	Student Services	ML; DL	
[106]	Demographic; Grade/GPA; Enrollment/Pathways; Protected Class	Grade Prediction	ML	
[111]	Demographic; Grade/GPA; Enrollment/Pathways; Protected Class	Retention	ML	Speculative Design
[113]	LMS/Engagement	Grade Prediction	Statistical Methods	
[116]	Demographic; Protected Class	Grade Prediction	ML; DL	
[117]	Demographic; Grade/GPA; Protected Class	Grade Prediction	DL	
[118]	Institutional; Enrollment/Pathways	Institutional Planning	ML;	