

Copyright is owned by the Author of the thesis. Permission is given for a copy to be downloaded by an individual for the purpose of research and private study only. The thesis may not be reproduced elsewhere without the permission of the Author.

Learning Analytics: On Effectiveness of Dashboarding for Enhancing Student Learning

A Thesis with Publications Presented in Partial Fulfilment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY
IN
INFORMATION TECHNOLOGY

School of Mathematical & Computational Sciences
Massey University, Auckland, New Zealand

Gomathy Ramaswami
November 2022

Abstract

Ongoing advancements in learning analytics have provided institutions with immense opportunities to identify and discern student learning patterns across different course offerings. These patterns can help identify those students who may be at some risk of course failure (or of course completion) as soon as possible, which further allows institutions in timely offering them guidance and support for overcoming their learning difficulties. Learning Analytics Dashboards (LAD) are currently used to deliver graphical representations of data-driven insights timeously to support management teams, instructors, and students. LADs provide a comprehensive overview on current learning environments with much use of visualizations for displaying learning patterns that can capture various aspects of the student learning experience. Hence, LADs are increasingly being used as a pedagogical approach for motivating students and supporting them in meeting their learning goals.

This research study has developed a student-facing LAD that shows a snapshot of students' online learning behaviors by implementing descriptive analytics components and also incorporates machine learning in a way that enables both predictive and prescriptive analytics. The study is divided into two parts. First, a generic predictive model has been developed to identify the at-risk students across a wide variety of courses. After generating the predictive model, model explainability using anchors has demonstrated the reasoning behind predictive models to enable transparency of the predictive models and increase students' trust as they interact with the LAD. Machine learning models used in this study have implemented prescriptive components by prioritizing which changes in learning behaviors and which learning strategies adopted by a student will most likely translate into favorable results.

Second, a LAD is developed. The dashboard provides visualizations that incorporates graphical and statistical information of online behavioral student patterns as they engage with the coursework. An online student survey that gauged LAD effectiveness for its usefulness and the

motivational impact of its prescriptive output to better engage students with the coursework has shown promising results. The LAD design, as far as we know, is the first in the learning analytics domain that has combined all three analytics, namely descriptive, predictive, and prescriptive. This thesis has investigated an active area of research and has paved the way for more meaningful LAD design and implementation, thereby contributing to both theory and practice.

Authors Declaration

This thesis was produced according to Massey University's "PhD thesis by publication" guidelines. This thesis is based on research that has been published in Journals and IEEE conference. In accordance with Journal and IEEE copyright policy, this thesis is containing all published manuscript. Consequently, there may be stylistic differences between the enclosed work or the published versions. Furthermore, the work contained within this thesis was published in several Journals. Each manuscript has addressed different problems and proposed efficient models to solve those problems. However, this may have some repetition in literature review.

I, Gomathy Ramaswami declare that this thesis is a fulfilment of the requirements for the fulfilment of the degree of Doctor of Philosophy (Ph.D.), from School of Mathematical & Computational Sciences, Massey University, is wholly my own work unless otherwise referenced acknowledged. This document has not been submitted for qualifications at any other academic institution.

Gomathy Ramaswami

Acknowledgements

First and foremost, I would like to express my sincere gratitude to my supervisory team: Dr. Teo Susnjak, and Dr. Anuradha Mathrani for their encouragement, direction, guidance, and support throughout my doctoral journey. Without their motivation, continuous encouragement and contributions, this research would not have been successfully completed. I am grateful to Andrew Rowart, Director, Online Learning Environment, Massey University, for providing me with the data needed for my research.

I sincerely acknowledge the kindness of Massey University GRS and the School of Natural and Computational Sciences for supporting my PhD studies with a Doctoral Scholarship and for providing financial support in the form of conference travel grants.

I would like to thank my fellow postgraduates for the friendship, support, and encouragement they provided.

A special thanks to my parents and my brother for their unconditional love and support.

Finally, I am fortunate enough to have a lovely son Sai Dhanwin whom I thank for his constant support and sacrifice which helped me in successfully completing the PhD journey on time.

Funding

I gratefully acknowledge the fundings received from:

- Massey Doctoral Scholarship provided by the Massey University.
- School of Natural and Computational Sciences (SNCS) for open access and conference registration fees.

Contents

Abstract	i
Authors Declaration	iii
Acknowledgements	iv
Fundings	v
List of Figures	xiv
List of Tables	xvi
1 Introduction	1
1.1 Introduction	2
1.2 Background of the research	4
1.2.1 Learning Analytics Dashboard Design	4
1.2.2 Target User and Purpose	4
1.2.3 Key Performance Indicators	5
1.3 This Study's context	6
1.4 Research objectives	7
1.5 Research Road Map.....	8
1.6 Contributions	10
1.6.1 Published Work	11
1.6.2 Submitted Articles	12
1.7 Thesis Outline	12

2 Use of Predictive Analytics within Learning Analytics Dashboards: A Review of Case Studies	15
2.1 Introduction	19
2.2 Overview of Systematic Literature Reviews on LADs	20
2.3 Research Questions	21
2.4 Methodology	21
2.5 Review of Dashboards Meeting the Inclusion Criteria	22
2.6 Results	29
2.6.1 Key Themes	29
2.6.2 Target Users	29
2.6.3 Types of Data Used for prediction	30
2.6.4 Prediction Methods	30
2.6.5 Evaluating the Prediction Accuracy	31
2.6.6 Visualization Methods	32
2.6.7 Dashboard Evaluations	32
2.7 Discussion	33
2.8 Recommendations for LAD Design	35
2.9 Conclusion	35
3. Predicting Students Final Academic Performance Using Feature Selection Approaches	40

3.1 Introduction	43
3.2 Existing Reviews	43
3.3 Datasets	44
3.3.1 Moodle Action Logs	44
3.3.2 Ethics Approval	44
3.3.3 Tools Used	44
3.4 Methodology	45
3.4.1 Model Evaluation Metrics	45
3.4.2 Feature Selection	45
3.5 Results and Discussion	46
3.5.1 Experiment 1	46
3.5.2 Experiment 2	46
3.6 Conclusion and Future Scope	47
4. Perspectives on the challenges of generalizability, transparency, and ethics in predictive learning analytics.....	48
4.1 Introduction	51
4.2 Research Questions	52
4.3 Current Studies on Learning Analytics	52
4.4 Challenges in Learning Analytics	53
4.4.1 Generalizability Challenges	53

4.4.2 Model Transparency (Interpretability and Explainability) Challenges	55
4.4.3 Ethical Challenges	55
4.5 Conclusions, Limitations, and future Directions	57
5. On Developing Generic Models for Predicting Student Outcomes in	
Educational Data Mining	60
5.1 Introduction	63
5.2 Review of Related Studies	64
5.2.1 Related Study Exploring Prediction Using Generalized Model	65
5.2.2 Related Study Exploring Early Prediction of Student Performance	65
5.3 Datasets	67
5.3.1 Action Logs from Moodle	67
5.3.2 Enrolment Management System	67
5.3.2 Student Management System	68
5.4 Method	68
5.4.1 Predictive Modelling	70
5.4.2 Machine Learning Procedure	71
5.4.3 Model Evaluation	71
5.4.4 Experimental Design.....	71
5.5 Results	71
5.5.1 Overall Performance Comparison of Classifiers	72

5.5.2 Classifier Performance Snapshots	72
5.5.3 Feature Importance	73
5.6 Discussion	75
5.7 Conclusions	76
6. Supporting Students' Academic Performance Using Explainable Machine	
Learning with Automated Prescriptive Analytics	79
6.1 Introduction	82
6.2 Related Work	83
6.2.1 eXplainable Artificial Intelligence (XAI).....	83
6.2.2 Post Hoc Explanations of Machine Learning Models.....	84
6.3 Analytics Workflow	84
6.3.1 Standard Predictive Analytics Workflow	85
6.3.2 Model Explainability	85
6.3.3 Prescriptive Analytics.....	86
6.4 Methods	86
6.4.1 Dataset and Features.....	86
6.4.2 Tools	87
6.4.3 Machine Learning Algorithms and Evaluation.....	87
6.4.4 Feature Importance.....	88
6.5 Result	90
6.5.1 Model Explainability Example using Anchors	90

6.5.2 Prescriptive Analytics Example using Counterfactuals	91
6.6 Discussion	93
6.7 Conclusions.....	94
7. Learning analytics dashboard: a tool for providing actionable insights to learners	96
7.1 Introduction	99
7.2 Analytics layers	100
7.3 Aims.....	102
7.4 Analytics in education	103
7.5 Challenges for building LADs	103
7.6 Representation and actions	104
7.7 Ethics	106
7.8 Agility.....	107
7.9 Methods	108
7.10 Analysis	110
7.11 Review of dashboards from literature	110
7.11.1 LADs with descriptive analytics capabilities.....	110
7.11.2 LADs with predictive analytics capabilities	110
7.11.3 LADs with prescriptive analytics capabilities	111
7.11.4 Reported effectiveness of LADs	111

7.12 Dashboard data analysis	112
7.13 Proposed learning analytics dashboard.....	113
7.14 Analytics Layers.....	114
7.15 Dashboard design details	115
7.16 Discussion	116
7.17 Future directions	117
7.18 Study limitations	118
7.19 Conclusion	118
7.20 Theoretical implications	118
7.21 Practical implications	119
8. Effectiveness of a Learning Analytics Dashboard for Increasing Student Engagement Levels	
8.1 Introduction.....	125
8.2 Background	126
8.2.1 Social comparison theory	126
8.2.2 Learning analytics dashboards	127
8.2.3 Predictive model explainability	127
8.2.4 LAD evaluations	127
8.3 Methods	128

8.3.1 Participants	128
8.3.2 Dashboard Evaluation	128
8.3.2.1 Statistical analysis	128
8.3.2.1 Survey analysis	129
8.4 Dashboard Analytics	129
8.4.2 Predictive Analytics	129
8.4.2 Model Interpretability	130
8.4.3 Model Explainability	131
8.4.4 Prescriptive analytics	131
8.5 Dashboard Architectural Layers	131
8.6 Dashboard Layout	132
8.7 Results of Dashboard Effects	134
8.7.1 Statistical Results	134
8.7.2 Survey Results	134
8.8 Discussion.....	135
8.7.1 Study Limitations	136
8.9 Conclusion.....	136
9. Conclusions, Limitations and Future Directions	
9.1 Conclusions.....	141
9.2 Limitations and Future Research Directions	144

10. Appendices	
Appendix 1: Capitalizing on Learning Analytics Dashboard for Maximizing Student Outcomes	151
Appendix 2: Ethics Approval forms	157

List of Figures

1.1 Roadmap of the Study	9
2.1 Methodology used for inclusion and exclusion	23
2.2 Trend in number of LAD studies as a three-year rolling average	27
2.3 Target users of the dashboards. Some dashboards target multiple audiences	29
2.4 Visualization types	31
2.5 Recommendations for Learning Analytics Dashboard Design.....	35
3.1 Frequency of Events (activities) for a course	44
3.2 Steps involved in prediction	45
3.3 Logistic Regression performance with and without feature selection method	46
4.1 Generalizability Challenges	55
4.2 Trade-offs between Low and High Model Transparency	56
4.3 Research Ethics Protocol.	57
5.1 Comparison of prior score with final score of students	69
5.2 F-measure and AUC plot of CatBoost on the Computer Applications and the Information Age—Semester 03 hold-out dataset... ..	73
5.3 Feature importance ranking plot with CatBoost SHAP.....	74
5.4 Shapley force plot of a sample.	74
5.5 Two SHAP scatter dependence plots.	75
6.1 Analytics workflow followed in this study.....	85

6.2 Predictive performances of various classifiers.....	88
6.3 Feature importance plot.	89
6.4 Two SHAP scatter dependence plots.....	89
6.5 (a) Anchor explanation for Student A. (b) Anchor explanation for Student B.....	91
6.6 (a) Counterfactual explanation translation for Student A. (b) Counterfactual explanation translation for Student B.....	93
7.1 Methodology used in this systematic review (Moher et al., 2009)	108
7.2 Total number of published articles presenting LADs that are covered in this study. The number of publications for 2021 is listed up to September of that year	109
7.3 Learning analytics dashboard designed for students	114
8.1 Feature importance plot with SHAP	131
8.2 LAD architecture	132
8.3 The student-facing LAD used for this study	133
8.4 Survey results	135

List of Tables

1.1 Category of indicators identified by Schwendimann et al. [21].....	5
2.1 Reviewed papers overviews	24
3.1 Information About the Course	44
3.2 Actions Performed by a Student in Moodle.....	45
3.3 Features Selected Using Embedded Method	46
3.4 Performance of Various Classifiers.....	46
5.1 Machine learning algorithms applied in educational environments.....	67
5.2 Information about the courses and corresponding log attributes	68
5.3 Feature description	70
5.4 Performance scores of various classifiers.....	72
6.1 Feature descriptions	86
6.2 Performance scores of various classifiers	87
6.3 a) Counterfactual explanation for Student A. (b) Counterfactual explanation for Student B	92
7.1 Reviewed papers overview.	101
7.2 Dashboard technologies and size of study cohorts.....	113
8.1 Feature description.	129
8.2 Performance scores of various classifiers	130

Chapter 1

Introduction

1.1 Introduction

Advancements in information learning technologies have facilitated the digitization of education and made education more accessible to students. Instructors can monitor students' learning processes by evaluating their learning portfolios using a learning management system (LMS) in an e-learning environment [1]. LMSs are web-based learning systems that utilize a virtual platform for enabling teaching and learning, such as offering the course content online, tracking student interactions, empowering peer communication over online forums, delivering course assessments to students, or releasing assessment grades [2]. Various stakeholders have different reasons for using an LMS. For instance, Romero and Ventura [3] suggest that students use an LMS to personalize their learning, such as reviewing specific material or engaging in relevant discussions as they prepare for their exams. Meanwhile, instructors rely on an LMS to deliver their course content and manage teaching resources in a relatively simple and uniform manner. Irrespective of how an LMS is used, user interaction with the system generates significant and detailed digital footprints that offer opportunities for advancing the field of learning analytics (LA) [4].

LA has the potential to interpret the data collected from the institutional sources to provide actionable intelligence about individual and groups of students for enhancing learning outcomes [5] and for building more context around how their learning environment is operating [6, 7]. This could involve predicting which coursework aspects are related to learners' success in the course, or providing relevant, personalized, and timely feedback to students and instructors [8, 9]. The simplest is descriptive analytics which uses current and historical data to identify past behaviors and learning trends from students' data. The most widespread use of LA-centered machine learning (ML) algorithms is in predictive analytics for analyzing students' engagement data with a course to develop models that can predict to identify students

most at risk of underperforming at an early stage [10]. This in turn can benefit students to further improve their academic performance [11, 12] and instructors can then support the at-risk students in overcoming their learning difficulties [13]. However, the most complex and possibly richest insight can be gained from prescriptive analytics which utilizes the underlying models and provides users with recommendations and suggestions about which particular behavioral changes are most likely to result in positive results.

However, majority of research into predicting academic performance has focused on developing tailor-made models which may likely perform poorly across numerous courses for reasons such as some courses may have small cohorts or newly constructed courses that do not have historic data from which patterns can be learned through machine learning. Such challenges which are a consequence of disconnect with historic data and current real-life data can cause poor model generalization [14]. Furthermore, offering model transparency into the mechanics of respective predictions is crucial for students to comprehend the algorithmic predictions so that they can act on the model's resulting predictions in an informed manner. In addition, by providing automated prescriptive feedback that is based upon underlying data-driven ML methods, we can advise students on what behavioral adjustments are more likely to result in more positive outcomes for them [15].

LA aims to assist students in achieving their learning goals by providing them with tools that gauge their learning process based on their online course interactions and create actionable feedback. Research has shown that, for instructors, access to tailored visual representations of relevant, up-to-date data, including constructive feedback, has the power to change how they teach and how they support their students to achieve better outcomes [16]. An effective way to present this data is via a dashboard which is a visual display that brings information together and guides the users in their decision-making processes. A LAD can utilize multiple data visualizations to present complex data [5] in a way in which it can be easily understood and

used by both instructors and learners. Benefits include identifying the at-risk students to provide them with customized assistance [9], increasing learner participation and engagement in courses [17], and generally supporting both educators and learners in enhancing the teaching and learning experiences [18, 19].

1.2 Background of the Research

Before developing a dashboard, it is important to understand the dashboard design, key performance indicators, visualization information, and the targeted users of the dashboard. This provides a perspective on the underlying reasons relating to the purpose of the dashboard development, that is, which design is appropriate, who are the target users and what are the key performance indicators that need to be addressed. The following subsections provide a short overview.

1.2.1 Learning Analytics Dashboard Design

A dashboard, as defined by Few [20] is "a visual display of the most important information needed to achieve one or more objectives; consolidated and arranged on a single screen so the information can be monitored at a glance" (p. 26). In learning analytics, a dashboard is specifically viewed as "a single display that aggregates one or multiple visualizations of different indicators about learner(s), learning process(es), and/or learning context" (p.37) [21, 22].

1.2.2 Target User and Purpose

Based on target users, the dashboard interface can be classified into four types: students, instructors, study advisers, and administrators [21]. From the perspective of purpose, Schwendimann et al. classified learning analytics dashboards as self-monitoring, monitoring others, and administrative monitoring and the primary target users of learning dashboards are teachers and learners. As the objectives of target users differ, the context and visual design of

the dashboard should also support the goals of the target users and allow the user to analyze and interpret available data easily based on individual interests [23].

Learning analytics dashboards designed from the learners' perspective often aim to enhance learning by allowing learners to set appropriate learning goals based on learning guidelines and give correct and customized feedback to enhance learning motivation, etc. In comparison, dashboards designed from instructors' perspectives often provide a more comprehensive overview of individual and group learning progress and seek to improve instructors' teaching strategies and support adaptive teaching.

1.2.3 Key Performance Indicators

Learning dashboards can also be viewed as a visual representation of key performance indicators [20]. Over 200 indicators were identified by Schwendimann et al. [21] from the literature and mapped these into six categories: learner-related, result-related, content-related, action-related, social-related, and context-related.

Indicator category	Description	Examples
Learner-related	Basic student information	Age, previous education, previous courses, competencies,
Result-related	Learning outcomes	Average grade, individual grade, class grade distribution
Content-related	Information about the content	Number of concepts, topics covered, links in the concept map
Action-related	Learning actions performed by students	Time spent on tasks, number of page visits, timeline
Social-related	How students interact with each other	Group discussions
Context-related indicators	About the learning environment	The learning location, placement in a classroom

Table: 1 Category of indicators identified by Schwendimann et al. [21].

However, presenting the complex data stored in learning systems more accessible, meaningful, and useful to users (i.e., instructors, students, student advisors and administrators) is a challenging task. But with a combination of data analytics and data visualization methods, we can make it possible for users to gain critical insights from learner patterns and delve into data

relationships that would otherwise have remained hidden [24]. A process proposed by Shneiderman [25] “Overview first, zoom and filter, then details-on-demand” has proven useful in guiding the design of interactive visual representations of statistical data (p. 2). Choosing the right tools, in this case, the most appropriate type of visual representation is crucial. Stoltzman [26] looked at several types of charts, regarding their suitability to visually represent specific actions or objectives. For trends over time, column graphs or line charts were suggested. Options for comparisons of data included area charts, bar charts, bullet graphs, column graphs, line charts, and scatter plots. Correlations could best be illustrated via line charts or scatter plots. For distributions, bar charts, box plots, or column graphs were recommended. Compositions could best be represented by doughnut charts, pie charts, stacked bar charts, or stacked column charts.

Specific intended goals and visualization types should be distinguished based on the perspective of the end user. In addition to supporting instructors to observe individual learning progress, an overview of class-level performance, for instance, the average performance of a group, a comparative overview, etc., can be helpful to support adaptive teaching [27]. Although the target groups may differ, the popularity of specific visual representations of the data in the learning analytics dashboard did not vary significantly [21]. The most popular visual representations are bar graphs, line graphs, tables, pie graphs, and network graphs, respectively [21].

1.3 This Study's Context

Increasing students’ motivation and helping them to self-reflect on their learning processes is a significant driver behind LA research [28]. This research study aims to develop a LAD that will enable students to monitor their engagement with their courses in real time. The underlying assumption is that in having access to real-time feedback on their learning engagement, it will

serve as a motivating factor and lead to positive modifications of student learning behavior which may result in more positive learning outcomes.

There are two technological components to the dashboard. The first is a data visualization tool that provides students with a snapshot of how they are tracking their levels of engagement and their comparative engagement levels with the rest of the class. The second component employs machine learning algorithms that extract learning and engagement patterns from the data of previous student cohorts to generate various models. For example, a model that can predict non-completion or unachieved learning outcomes early in a course, based on which students exhibit similar engagement patterns to those of previous students (from earlier cohorts).

1.4 Research Objectives

The study's objectives are two fold. The first is to develop various analytics technologies namely descriptive, predictive, and prescriptive technologies to support the LAD. In doing so, this research will seek to answer the following questions:

1. Determine how a generic, course-agnostic predictive model can be developed to predict students' outcomes across disparate courses.
2. How can explainable ML approaches be leveraged to provide effective and human-understandable reasoning behind their conclusions about a student's academic performance?
3. How can automatic data-driven feedback and actionable recommendations be derived?

Next, build a dashboard plug-in by integrating the technologies and deploy the dashboard extension for some ongoing courses. User learning traces captured from the dashboard will be evaluated to study the effectiveness of the dashboarding strategy to answer the following research questions:

1. Determine how the identified model can be projected to students so that the students can view their current academic performance as putting them at low risk, or high risk of not meeting their learning outcomes.
2. Does the student-facing analytical dashboard result in more student engagement with the institutional LMS?
3. What are the students' perceptions of the dashboard's effect on their learning performance?

1.5 Research Roadmap

The following figure illustrates the roadmap of the study which was followed to answer the posed research questions and achieve study objectives.

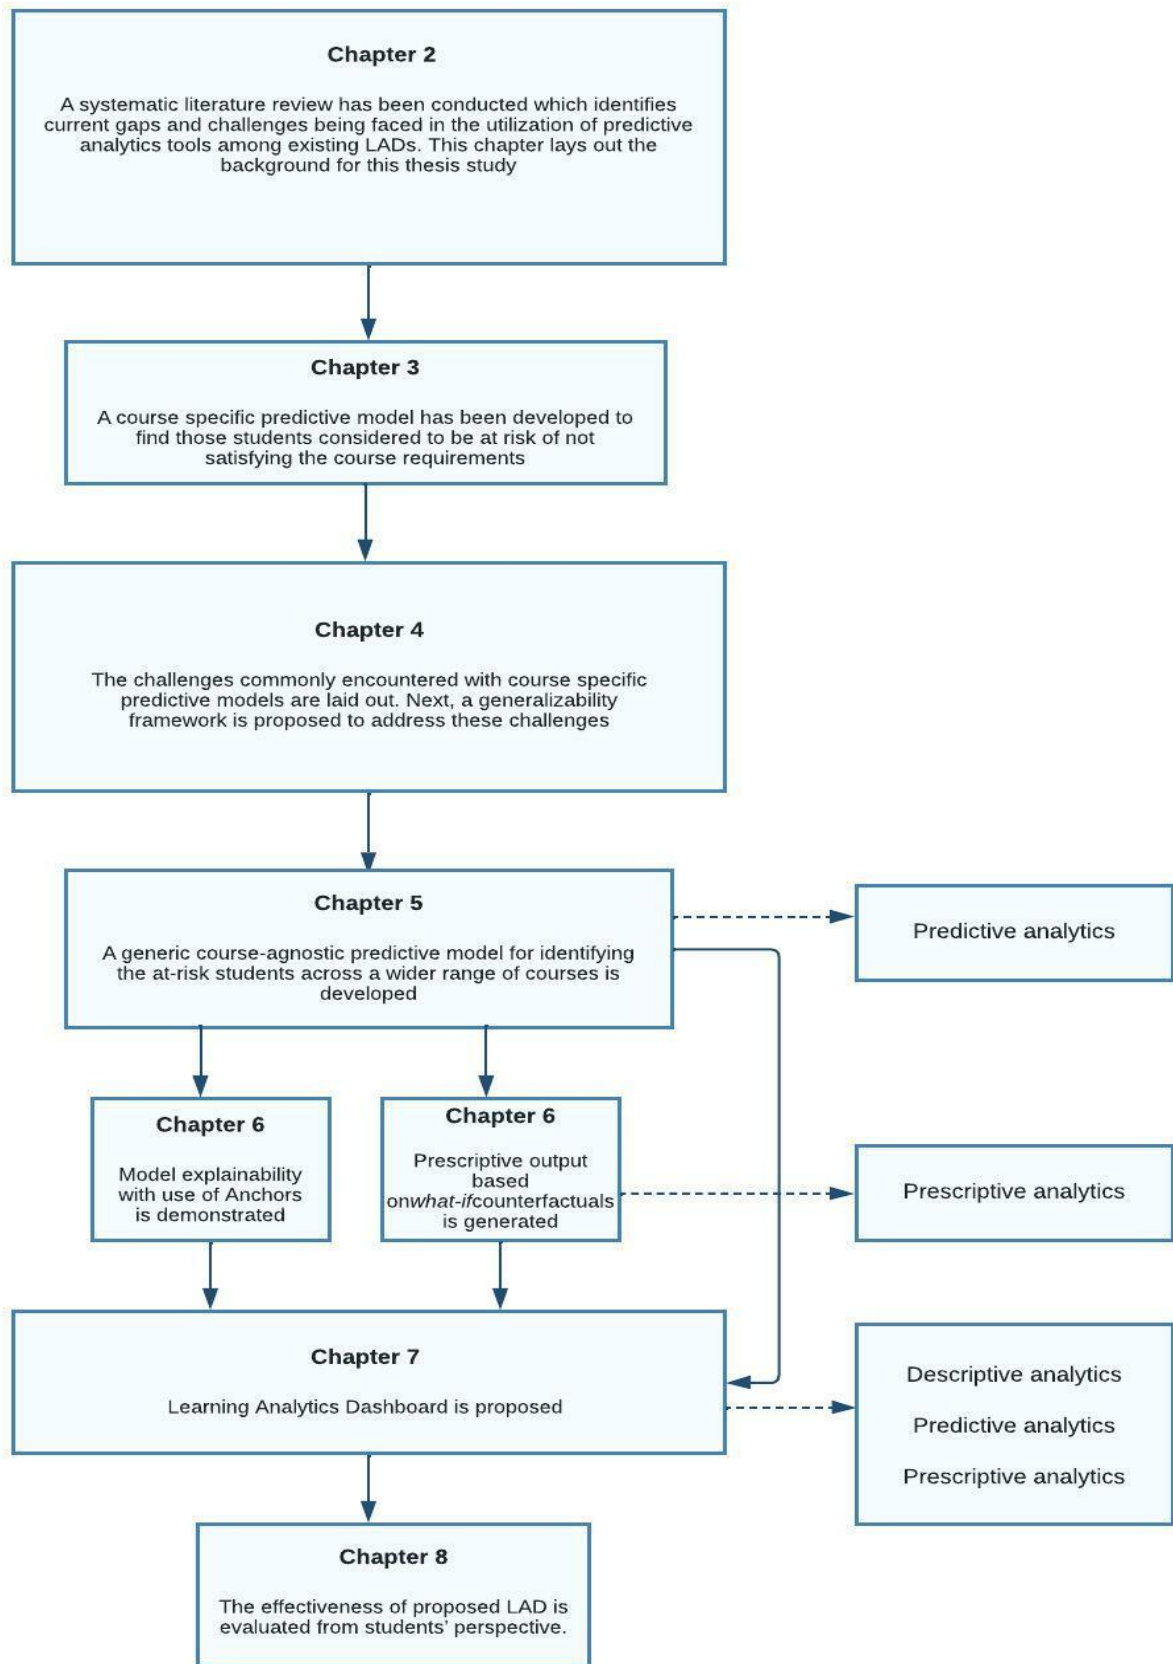


Figure 1. Roadmap of the study

1.6 Contributions

The main contributions of the thesis are as follows:

This study demonstrates a detailed study about the portability of predictive model between university courses that helps to identify at-risk students. This study is a stepping stone to move to the stage of developing multiple models across diverse pattern of courses that resulting in overfitting. The finding of this study has broader implications for institutions seeking generalized and portable models for identifying students at risk of academic failure.

To date, the major focus in LA has been with descriptive and predictive analytics. Nevertheless, prescriptive analytics are now seen as a future course of action. We refined these challenges to LAD projects and have identified the lack of agility in higher education institutions to reveal either the predictive decision or offer suggestion for students to perform better in their course. Taking these two gaps into account, this study has demonstrated how interpretability of predictive models can be made available to the learners and critically, how the specific predictions for a given learner can be explained to them. In doing so, we show how prescriptive analytics can be integrated in the model to generate data-driven advice that learners can relate to; thus, bringing about effectual learning changes for increasing the probability of realizing successful outcomes.

Finally, this study proposes a state-of-the art dashboard that not only leverages descriptive analytics components, but also integrates machine learning in a way that enables both predictive and prescriptive analytics. This study demonstrate how emerging analytics tools can be used to enable learners to adequately interpret the predictive model behavior, and more specifically to understand how a predictive model arrives at a given prediction. Furthermore, this study show how data-driven prescriptive analytics can be deployed within dashboards to provide concrete advice to the learners, and thereby increase the likelihood of triggering

behavioral changes. Additionally, our research establishes pathways for the institutions towards integrating all forms of analytics accessible to the wider audience.

The proposed dashboard is the first of its kind in terms of breadth of analytics that integrates three different analytics paradigms.

1.6.1 Published Work

The author of this work has contributed to the following published articles in peer-reviewed journals and conferences. These articles are presented next (in the order of their year of publication).

- **Ramaswami, G. S.**, Susnjak, T., & Mathrani, A. (2019, December). Capitalizing on Learning Analytics Dashboard for Maximizing Student Outcomes. In *2019 IEEE Asia-Pacific Conference on Computer Science and Data Engineering (CSDE)*, 1- 6. IEEE. doi:10.1109/csde48274.2019.9162357
- **Ramaswami, G. S.**, Susnjak, T., Mathrani, A., & Umer, R. (2020, December). Predicting Students Final Academic Performance using Feature Selection Approaches. In *2020 IEEE Asia-Pacific Conference on Computer Science and Data Engineering (CSDE)*, 1-5. IEEE. doi:10.1109/csde50874.2020.9411605.
- Mathrani, A., Susnjak, T., **Ramaswami, G.**, & Barczak, A. (2021). Perspectives on the challenges of generalizability, transparency, and ethics in predictive learning analytics. *Computers and Education Open*, 2, 1-9. <https://doi.org/10.1016/j.caeo.2021.100060>
- Susnjak, T., **Ramaswami, G. S.**, & Mathrani, A. (2022). Learning analytics dashboard: a tool for providing actionable insights to learners. *International Journal of Educational Technology in Higher Education*, 19(1), 1-23. <https://doi.org/10.1186/s41239-021-00313-7>

- **Ramaswami, G.,** Susnjak, T., & Mathrani, A. (2022). On Developing Generic Models for Predicting Student Outcomes in Educational Data Mining. *Big Data and Cognitive Computing*, 6(1), 6, 1-16. <https://doi.org/10.3390/bdcc6010006>
- **Ramaswami, G.,** Susnjak, T., Mathrani, A., & Umer, R. (2022). Use of Predictive Analytics within Learning Analytics Dashboards: A Review of Case Studies. *Technology, Knowledge, and Learning*, 1-22. <https://doi.org/10.1007/s10758-022-09613-x>
- **Ramaswami, G.,** Susnjak, T., & Mathrani, A. (2022). Supporting Students' Academic Performance Using Explainable Machine Learning with Automated Prescriptive Analytics. *Big Data and Cognitive Computing*, 6(4), 105, 1-14. <https://doi.org/10.3390/bdcc6040105>

1.6.2 Article in press

Ramaswami, G., Susnjak, T., & Mathrani, A. Effectiveness of a Learning Analytics Dashboard for Increasing Student Engagement Levels. *Journal of learning analytics*.

1.7 Thesis Outline

The thesis is divided into eight chapters. This chapter discusses the background of the study, presents the objectives and research questions, and the scope of the study. The contributions from this study are stated.

Chapter 2 provides a systematic literature review to identify gaps and challenges in the current research into the utilization of predictive analytics tools among existing LADs, which consequently forges a path for future research. The systematic literature review presented below covers the methodology of the review, the results of the review, an overview of the data, and an overview of machine/data mining methods used for the prediction of student performance.

Chapter 3 presents an analysis of LMS data to see how accurately student academic performance can be forecasted when their weekly engagement data is integrated with assignment scores. This study highlights the selection, using feature selection methods, of features to enable early prediction of student academic performance.

Chapter 4 presents the challenges encountered in building individual predictive models for every single course and provides a generalisability framework to address them.

Chapter 5 demonstrates how a generic predictive model can be developed that identifies at-risk students across a wide variety of courses.

It was found from the existing studies in the educational domain that there was a lack of model interpretability and explainability of outputs which may lower the utility, and over time grind down the trust of users. Hence, Chapter 6 provides findings based on model interpretability and explainability of their predictions to students, and with counterfactuals which explicitly demonstrate alternate outcomes for the student if a behavioral change were to take place in specific areas.

Chapter 7 conducts a systematic literature review to identify the strengths and weaknesses among the existing dashboards and propose a LAD that brings descriptive, predictive, and prescriptive analytics into one display.

Chapter 8 conducts the impact study of the developed LAD to evaluate whether dashboard access has encouraged students to engage more with the learning management system (LMS) and also find out the students' perceptions of the dashboard's effect on their learning performance.

Chapter 9 summarizes the study, discuss the limitation of the study, and provides future research direction.

Chapter 2

Use of Predictive Analytics within Learning Analytics Dashboards: A Review of Case Studies

The contents of this chapter are from the following article. In accordance with the SpringerLink open access policy, any full text that has been included, is the unmodified accepted article.

©2022 SpringerLink. Reprinted, with permission, from: Ramaswami, G., Susnjak, T., Mathrani, A., & Umer, R. (2022). Use of Predictive Analytics within Learning Analytics Dashboards: A Review of Case Studies. *Technology, Knowledge and Learning*, 1-22.

<https://doi.org/10.1007/s10758-022-09613-x>

The SpringerLink open access policy, permits the use, sharing, adaptation, distribution and reproduction in any medium or format as long as appropriate credit goes to the authors.

SpringerLink does not endorse any of Massey University's products or services. Internal or personal use of this material is permitted. If interested in reprinting/republishing Springerlink copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribution, please go to:

<https://journalofbigdata.springeropen.com/submission-guidelines/> copyright to learn how to obtain a license from the correct link.

STATEMENT OF CONTRIBUTION DOCTORATE WITH PUBLICATIONS/MANUSCRIPTS

We, the candidate and the candidate's Primary Supervisor, certify that all co-authors have consented to their work being included in the thesis and they have accepted the candidate's contribution as indicated below in the *Statement of Originality*.

Name of candidate:	Gomathy Suganya Ramaswami
Name/title of Primary Supervisor:	Dr. Teo Susnjak
In which chapter is the manuscript /published work:	Chapter 2
Please select one of the following three options:	
<input checked="" type="radio"/> The manuscript/published work is published or in press <ul style="list-style-type: none"> • Please provide the full reference of the Research Output: Ramaswami, G., Susnjak, T., Mathrani, A., & Umer, R. (2022). Use of Predictive Analytics within Learning Analytics Dashboards: A Review of Case Studies. <i>Technology, Knowledge and Learning</i>, 1-22. https://doi.org/10.1007/s10758-022-09613-x 	
<input type="radio"/> The manuscript is currently under review for publication – please indicate: <ul style="list-style-type: none"> • The name of the journal: • The percentage of the manuscript/published work that was contributed by the candidate: • Describe the contribution that the candidate has made to the manuscript/published work: 	
<input type="radio"/> It is intended that the manuscript will be published, but it has not yet been submitted to a journal	
Candidate's Signature:	Gomathy Ramaswami <small>Digitally signed by Gomathy Ramaswami Date: 2022.11.07 07:42:59 +13'00'</small>
Date:	07-Nov-2022
Primary Supervisor's Signature:	Teo Susnjak <small>Digitally signed by Teo Susnjak Date: 2022.11.07 15:30:10 +13'00'</small>
Date:	7-Nov-2022

This form should appear at the end of each thesis chapter/section/appendix submitted as a manuscript/publication or collected as an appendix at the end of the thesis.



Use of Predictive Analytics within Learning Analytics Dashboards: A Review of Case Studies

Gomathy Ramaswami¹ · Teo Susnjak¹ · Anuradha Mathrani¹ · Rahila Umer²

Accepted: 30 June 2022 / Published online: 26 August 2022
© The Author(s) 2022

Abstract

Learning analytics dashboards (LADs) provide educators and students with a comprehensive snapshot of the learning domain. Visualizations showcasing student learning behavioral patterns can help students gain greater self-awareness of their learning progression, and at the same time assist educators in identifying those students who may be facing learning difficulties. While LADs have gained popularity, existing LADs are still far behind when it comes to employing predictive analytics into their designs. Our systematic literature review has revealed limitations in the utilization of predictive analytics tools among existing LADs. We find that studies leveraging predictive analytics only go as far as identifying the at-risk students and do not employ model interpretation or explainability capabilities. This limits the ability of LADs to offer data-driven prescriptive advice to students that can offer them guidance on appropriate learning adjustments. Further, published studies have mostly described LADs that are still at prototype stages; hence, robust evaluations of how LADs affect student outcomes have not yet been conducted. The evaluations until now are limited to LAD functionalities and usability rather than their effectiveness as a pedagogical treatment. We conclude by making recommendations for the design of advanced dashboards that more fully take advantage of machine learning technologies, while using suitable visualizations to project only relevant information. Finally, we stress the importance of developing dashboards that are ultimately evaluated for their effectiveness.

Keywords Early warning system · Student feedback system · Systematic review · Learning analytics dashboard

✉ Gomathy Ramaswami
g.ramaswami@massey.ac.nz

Teo Susnjak
T.Susnjak@massey.ac.nz

Anuradha Mathrani
A.S.Mathrani@massey.ac.nz

Rahila Umer
rumer@buitms.edu.pk

¹ School of Mathematical and Computational Sciences, Massey University, Auckland, New Zealand

² Engineering and Management Sciences, Balochistan University of Information Technology, Quetta, Pakistan

1 Introduction

Learning analytics research focuses on the usage of learners' data for developing personalized applications that support a wide range of stakeholders (e.g., learners, instructors, advisers, administrators). Learners' data are analyzed for key performance indicators capturing learner behavior patterns, which are then visualized (Baker, 2016). The move towards visual displays has led to the development of learning analytics dashboards (LADs) which can prompt user reflection via relevant depicted insights, and potentially inform stakeholders on required interventions for optimizing current learning environments. Schwendimann et al. (2017, p. 37) defined LADs as "displays that aggregate different indicators about learner(s), learning process(es) and/or learning context(s) into one or multiple visualizations", which can also incorporate text-based content (Podgorelec and Kuhar, 2021). Though the layouts of individual LADs may differ, up to now, they primarily leverage two kinds of analytics, namely, descriptive and predictive (Afzaal et al., 2021). Descriptive analytics tends to look backwards to capture trends that are drawn from digital footprints left behind by learners while interacting within virtual environments. Digital footprints can result from accessing various learning resources or by undertaking collaborative activities or other communication interactions (Delen & Demirkan, 2013). Descriptive analytics also provides learners with snapshots of their current learning status, drawn again by tracking various online movements with the goal of empowering learners with self-knowledge about their learning behavior and learning progression at different points in time. In contrast, predictive analytics is forward looking. It provides forecasts about individual learners such as their probable performance in upcoming assignments and future grades on completion of a course. The integration of both predictive with descriptive analytics brings out richer insights, ultimately adding more value to LADs. However, as it currently stands, predictive analytics are only now beginning to emerge within LADs (Susnjak et al., 2022).

LADs serve multiple stakeholders. For one, they provide a scalable way for instructors and administrators to monitor student engagement in real-time. Emergent student activity patterns reveal actionable and useful insights, especially when instructors and students are physically separated from each other as in online environments. Students acquire a greater degree of self-knowledge through LADs by gaining more visibility into their online learning behaviors (Verbert et al., 2013a, b) which facilitate more informed study-related decisions. Most importantly, personalized learner-centric metrics on LADs can enable a form of self-reflection that results in positive behavioral adjustments. Meanwhile, instructor-facing LADs can identify at-risk students especially through predictive analytics, enabling instructors to initiate various forms of interventions which increase the probability of attaining successful course outcomes (Greller & Drachsler, 2012; Yoo et al., 2015).

Online learning environments are increasingly using broader types of dashboards to track student activities (Verbert et al., 2013a, b). These dashboards differ widely in the type of data used for analysis, the visual information displayed, the audience that they target, their overarching purpose (or theme) of the LADs as well as in their strategies to evaluate their effectiveness. To date, limited systematic literature reviews have been conducted summarizing the above aspects of LADs. Indeed, with the emergence of predictive analytics, no systematic literature review exists which considers the LADs from the perspective of what is being predicted, what input data is being used, which algorithms are employed and how the accuracy is evaluated. This review article seeks to fill this gap.

We first present an overview of the published literature reviews to build the groundwork already established in the LAD domain. The overview identifies the current state-of-the art and gaps in LADs, leading to three research questions to guide our study. Next, we present our own systematic literature review (SLR) in which we carry out an in-depth investigation into recent LAD case studies (published between 2011 and 2021) to answer these questions and provide recommendations for future LAD designs.

2 Overview of Systematic Literature Reviews on LADs

This section provides an overview of key findings covering the topic of LADs from four systematic literature reviews that were published between 2017 and 2019. These reviews were selected on the basis that they followed a structured article search process to identify ongoing trends in student-facing LADs and to assess the current state of this field.

Schwendimann et al. (2017) examined LADs based on categorization of learning contexts, data sources, visualizations and analysis types. Their review comprised 55 journal papers published between 2010 to 2015. They found that predominant theme of focus of the dashboards was on monitoring student progress, while being both student-facing and instructor-facing, and mostly relying on a single data source (i.e., log files from a learning management system (LMS)). Given that most studies were exploratory or proof-of-concept in nature, the conclusion was that very little analysis was performed on the effectiveness of LADs to impact learning outcomes.

Bodily and Verbert (2017) conducted an SLR on LAD studies published between 2005 and 2016. A total of 93 papers were reviewed from the perspective of LAD functionality, data sources used, design features as well as perceived and actual effects they elicit. The findings were inconclusive in respect to the effectiveness of LADs overall and specific designs, with the authors suggesting further research be undertaken.

Matcha et al. (2020) evaluated the impact of LADs on learning and teaching. From a review of 29 papers published between 2010 and 2017, the authors noted that existing LADs are insufficiently grounded in learning theory and do not offer insights into what constitutes effective learning approaches. The review suggests that evaluation of LADs should be done in learning contexts across numerous iterations for a stronger methodological foundation.

A recent systematic review conducted by Valle et al. (2021) focused on reviewing articles between 2012 and 2019 to determine whether evaluations of learner-facing LADs actually measured their efficacy. Similar to other reviews, a lack of alignment between the intended outcome and their evaluation measures was noted. The review added that current LADs are rarely designed to support students' self-regulated learning, and the study again recommended the need for appropriate measures for evaluating the efficacy of LADs.

All four LAD reviews focused on who the target users were, what types of data were used, and which evaluation method was used. Bodily and Verbert (2017) expanded the contribution by devising a LAD evaluation criterion that considers their effectiveness in how they impact learner behaviors and overall achievements. In their subsequent review, they also took account of the sample size of the participants and whether they were from a single or from multiple courses, while Schwendimann et al. (2017) differed in their perspective and categorized studies in terms of technologies used for presenting LADs to users.

Existing reviews have provided substantial contributions to LAD research; however, with the emergence of predictive analytics and its growing importance in enhancing LAD

capabilities (Bodily & Verbert, 2017; Susnjak et al., 2022), there is now a gap in the body of literature exploring this recent technological development. This is especially pertinent as predictive modelling is now recognized as an effective approach for identifying at-risk students, while also having the potential to improve both learner outcomes and retention rates (Umer et al., 2021). Hence, the point of difference in our systematic review is in examining LADs which incorporate predictive analytics. Accordingly, our research questions seek to address the previously unexamined role of predictive modelling within LADs, together with replicating and extending some findings from prior reviews.

3 Research Questions

RQ1 What are the key themes underlying LAD usage in existing literature? Further, who are the target users and what visualization techniques are employed?

RQ2 What types of data and which machine learning algorithms are commonly used for predictive analytics? What is being predicted and how is the prediction accuracy being assessed?

RQ3 What evaluation methods have been conducted for studying the effectiveness of LADs?

4 Methodology

Our systematic literature review followed guidelines proposed by Kitchenham and Charters (2007). The field of learning analytics (LA) and educational data mining (EDM) has grown significantly since 2010 (Park & Jo, 2019), hence this study examined literature published in journals and conferences between 2011 to 2021. This review lays exclusive emphasis on LA systems that collect student data, apply predictive modelling, while delivering the prediction results to students, instructors or academic advisers in the form of visualizations or text-based feedback. We included studies which communicated results either via dashboard displays or through reports sent as email attachments. Conferences provide a publishing outlet for emergent fields; accordingly, those conferences with an explicit focus on LA and EDM (e.g., International Conference on EDM and International Conference on LA and Knowledge) that met the inclusion criteria were considered.

Following a similar methodology used in earlier studies (Matcha et al., 2020; Schwendimann et al., 2017), we considered five main academic databases, namely, ACM Digital Library, Science Direct, IEEE Xplore, SpringerLink and Wiley. Additional databases including Google Scholar and Scopus were included in our SLR, as is recommended by Kitchenham and Charters (2007) for detecting relevant articles that could be overlooked since these databases are not typically indexed in common literature databases. Our search query comprised keywords pertinent to LAD literature (i.e., ‘learning analytics dashboard’ AND (‘students feedback system’ OR ‘early warning system’ OR ‘predictive analytics’ OR ‘visualization tool’)). This yielded 4350 articles across various academic data sources. Next, we removed duplicate papers and reduced to 2840. The exclusion criteria were applied where papers that were not written in English or contained less than 3 pages or were dissertations were removed. The papers that used MOOC datasets were also

filtered. We assessed titles and abstracts of each article and retained those studies which were related to LADs and early warning systems (EWS) with embedded predictive analytics. EWS studies were included since some used visual tools for delivering information. Finally, 11 journal articles and 4 conference proceedings (indexed in the Scopus/ACM digital library) were considered admissible for answering the research questions. Figure 1 describes the overall methodology.

5 Review of Dashboards Meeting the Inclusion Criteria

This section provides a summary of 15 LADs in view of the three research questions. Table 1 presents an overview of the reviewed papers highlighting key LAD characteristics prevalent in current LADs. Overall, we observed an increasing trend in the number of LAD studies which have been conducted over the search period. This can be seen in Fig. 2 which shows frequencies presented as a three-year rolling average in order to smoothen the data and remove noise.

Arnold and Pistilli (2012) from Purdue University created Course Signals, a student- and instructor-facing dashboard which tracks student progress. The data comprised student course performance, student engagement in terms of interactions with the LMS, academic history, and demographics. A visualization leveraging colors in the form of traffic lights to convey three risk levels, high (red), medium (yellow) and low (green), was developed and made available to the target users as early as the second week of the semester. However, the tool did not provide direct insights into model reasoning, that is, why a particular student was considered to be at risk of failure, making it difficult to recommend a specific remediation.

Essa and Ayad (2012) developed a student success system (S3) to alert instructors on student risk levels and to provide feedback to the corresponding students. Similar to Course Signals, S3 employed color as the primary visual encoding channel¹ in the form of a traffic light visualization to express risk probabilities. The authors used an ensemble modelling strategy which combined the predictions of several base models. S3 had no student-facing functionality, and no evaluation of the tool was conducted.

Agnihotri and Ott (2014) from the New York Institute of Technology designed a dashboard to identify at-risk students needing support, aiming to increase retention of first-year students. Survey, financial and pre-enrolment data were used for the tool. It was observed that students predicted to be at-risk by the model indeed aligned with those who did not return in the subsequent year.

Hu et al. (2014) developed an EWS to identify at-risk students. Algorithms such as C4.5, regression trees, logistic regression (LR), and boosting were applied to develop an EWS for online undergraduate courses. Learner data comprising login behavior, online course material interaction, assignment submission status and forum discussion activities were collected during three different time intervals in a semester. The instructors used the tool to adjust their teaching methods for poor-performing students, with the study concluding that EWS improved learners' performances and reduced attrition rates.

Jayaprakash et al. (2014) designed the Open Academic Analytics Initiative (OAAI) with an EWS aimed at identifying at-risk students. The authors claimed that the system

¹ Also referred to as a perceptual channel.

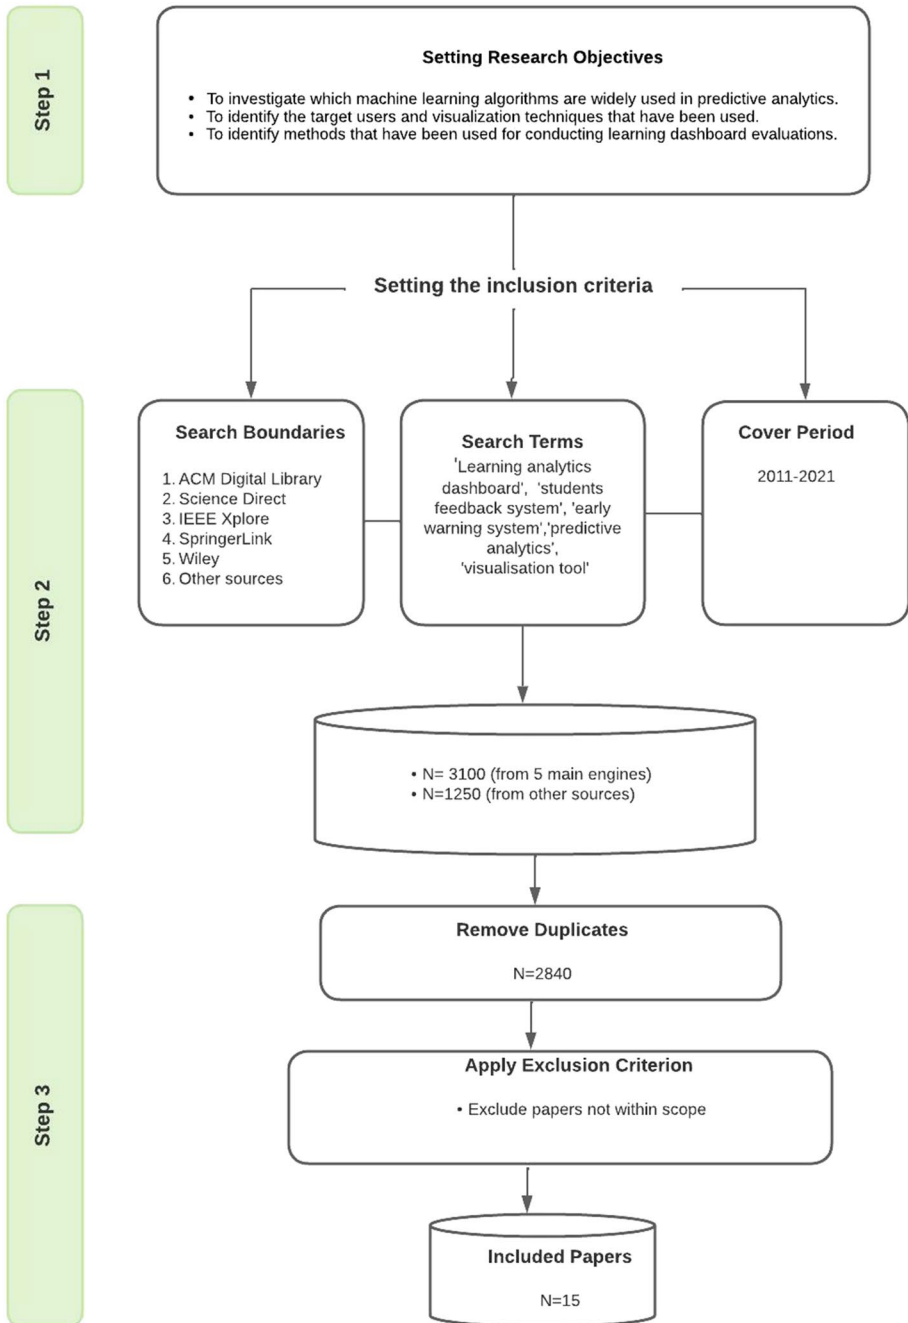


Fig. 1 Methodology used for inclusion and exclusion

Table 1 Reviewed papers overview

Key theme	Tool	Target users	Data used for prediction	Prediction target	Predictive algorithms	Prediction evaluation metric	Visualization Technique	Dashboard Evaluation Technique
Study Plan	Course Signal (Arnold & Pistilli, 2012)	Learners	Assessment grades; demographic data; pre-academic data; LMS data	Student risk level	Proprietary algorithm (i.e., student success algorithm)		Color-based custom charts	Mixed methods, and Impact evaluation conducted
	NYIT (Agnihotri & Ott, 2014)	Academic advisers	Enrolment data; and pre-academic data	Student risk level	Ensemble, Bayes, LR, DT, NN	Recall and precision	Text table	No evaluation conducted
	EWS (Hu et al., 2014)	Instructors and learners	LMS data	Student risk level	CART	Accuracy, FPR, FNR	Line chart; box-and-whisker plot; text table; pie chart (gauge plot)	Qualitative study conducted and impact evaluation not conducted
	OAAI (Jayaprakash et al., 2014)	Instructors and Learners	Demographic data; pre-academic data; and assessment grades	Student risk level	LR	Accuracy, ROC curves, precision, Recall, FP rate	No visualization provided, only email was sent to learners	Only impact evaluation conducted
	PredictED (Corrigan et al., 2015)	Learners	LMS data	Student risk level	SVM	ROC curves	No visualization provided only email was sent to learners	No evaluation conducted
	OU analyse (Kuzilek et al., 2015)	Academic advisers and instructors	Demographic, and LMS data	Risk level of successful assignment submission	Naïve Bayes, kNN and CART	Precision, Recall, F-measure	Bar chart; text table; color-based custom charts	Mixed methods conducted
	LISSA (Charleer et al., 2018)	Academic advisers	Prior grades; and assessment grades	Duration until course completion	Predictive algorithm was not used		Bar chart (standard and an histogram); color-based custom charts	Mixed methods conducted, and impact evaluation not conducted

Table 1 (continued)

Key theme	Tool	Target users	Data used for prediction	Prediction target	Predictive algorithms	Prediction evaluation metric	Visualization Technique	Dashboard Evaluation Technique
Monitoring Learning Progress	Prescriptive learning dashboard (He et al., 2018)	Academic advisers	Demographic; pre-academic data	Student risk level	RF and ensemble methods	ROC curves	Line chart; bar chart	Evaluation strategy is not mentioned
	Early warning dashboard (Bañeres et al., 2020)	Instructors and learners	LMS data, enrolment, and assessment grades	Student risk level	NB, DT, KNN, SVM	Accuracy, F score, True negative rate, True positive rate	Color-based custom charts; text table	Impact evaluation conducted
	LADA (Gutiérrez et al., 2020)	Academic advisers	Assessment grades, courses booked; academic history; demographic data	Student risk level	Clustering techniques (Fuzzy C-means algorithm)	Brier score	Bar chart (standard and slider); text table; pie chart (donut chart); color-based custom charts; scatter plot; radar chart;	Mixed methods conducted and impact evaluation not conducted
Monitoring Learning Progress	EWS (Plak et al., 2021)	Academic advisers and instructors	Demographic data; pre-academic data number of enrolled programmes	Student risk level	LR, RF, SVM	MAE	Text table	Mixed methods
	S3 (Essa & Ayad, 2012)	Instructors	LMS data	Student risk level	Predictive ensemble method	Accuracy	Sociogram, bar chart; scatter plot; color-based custom chart (risk quadrant)	Not mentioned

Table 1 (continued)

Key theme	Tool	Target users	Data used for prediction	Prediction target	Predictive algorithms	Prediction evaluation metric	Visualization Technique	Dashboard Evaluation Technique
	Academic EWS (Wang et al., 2018)	Learners	Library data; number of courses taken; attendance; and assessment grades	Student risk level	DT, Bayesian, and artificial neural network	Accuracy	Not mentioned	No evaluation conducted
	Student facing LAD (de Quincey et al., 2019)	Learners	LMS data; assessment grades	Student risk level	DT		Line chart; bar chart	Mixed methods conducted
	(Hellings & Haelermans, 2020)	Learners	Assessment grades; pre-academic data; LMS data; demographic data	Student risk level	Decision stump with AdaBoost		Pie chart (gauge plot); bar chart;	Impact evaluation conducted

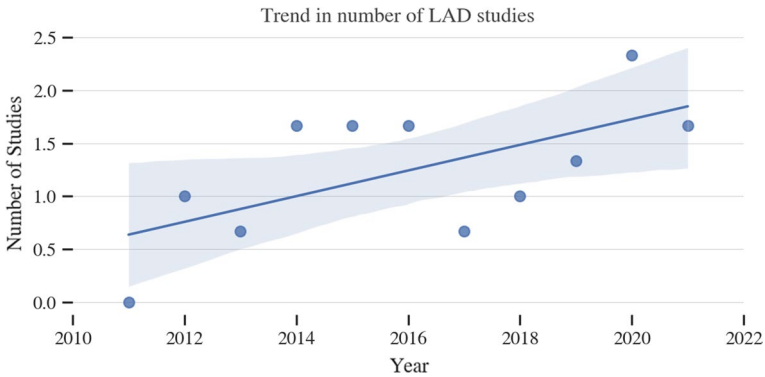


Fig. 2 Trend in number of LAD studies as a three-year rolling average

achieved accuracy of 84%. The system provided two intervention strategies, one entailed a message with guidance being sent to at-risk students, while the other involving online academic support. The study reported that students who received alert messages experienced a positive impact on their final outcomes although no clear gains in course grades were observed by the group receiving additional online academic support.

Corrigan et al. (2015) from Dublin City University developed PredictED to forecast students' final exam grades using student behavior data from Moodle (LMS). Support Vector Machines (SVM) was used to classify students as either passing/failing for each week of the semester until the exams. The opt-in students received weekly emails based on prediction results and showed an improvement in final grades of nearly 3% on average. Like Course Signals, this tool did not provide explainability of the model's reasoning making it difficult to recommend a specific remediation for at-risk students.

In another study, Kuzilek et al. (2015) from Open University developed a dashboard to predict at-risk students at early stages of their study. A dashboard with two views – a course view and a student view – was implemented. The course view provided an aggregated view of student activities and overall assessment results, while the student view presented a table with individual student results and predictions regarding the submission of the upcoming assessment. Weekly emails were sent to course coordinators. Demographic data and interaction data from the virtual learning environment (LMS) were used for the analysis. However, the tool did not provide any direct insight to students.

The goal of LISSA (Learning dashboard for Insights and Support during Study Advice) was to assist academic advisers in helping students plan a more attainable study programme and not waste time on unsuitable choices (Charleer et al., 2018). Student grades and historical data were used for creating a prediction bar depicting the duration of their progress in the bachelor programme. A custom color-based visualization was used that communicated predicted outcome categories of success, mediocre pass and fail as green, orange and red respectively. The tool was only used to support the advising session and did not target students directly.

A prototype LAD was developed by He et al. (2018) aiming to provide students with academic performance guidance for STEM courses. The authors used RF (Random Forest) for predicting students' results. A gauge plot presented the predicted results to student advisers who relied on the tool to initiate interventions. The dashboard did not

depict students' learning progression,. Neither was it deployed on real courses nor was any tool evaluation conducted.

Wang et al. (2018) designed an EWS with the goal of reducing student dropout and reducing graduation delays. This system included data from the library, dormitory, grades, attendance, and engagement. The library and dormitory data enabled closer monitoring of study habits. The NB (Naïve Bayes) algorithm attained classification accuracy of 86%, with grades and library data being among the key variables. The tool was neither productionized nor subjected to evaluation.

To motivate and personalize the student learning, de Quincey et al. (2019) developed a student-facing LAD which focused on showing student weekly progress along with their predicted semester-end scores. Similar to Hellings and Haelermans (2020), a weekly reminder email was sent to students during the first few weeks, with personalized recommendations also displayed on the dashboard. A mixed method LAD evaluation showed favorable feedback regarding only the usability of the dashboard.

Bañeres et al. (2020) presented two EWSs, one instructor-facing and the other student-facing, to identify at-risk learners and provide them with personalized feedback. A different model was built for each course based on students' grades using NB, Decision Tree (DT), K-Nearest Neighbour (kNN) and SVM algorithms. The authors found NB to be the best for their dataset; however, changes in course assessments led to inconsistencies in the predictions.

Gutiérrez et al. (2020) developed LADA (Learning Analytics Dashboard for Advisers) to support academic advisers in setting up students' semester plans. Data used in the analysis included student grades, course credits and lists of courses taken or enrolled into by students. The dashboard provided predictions on students' probability of success and a further indication regarding the confidence of predictions. An evaluation of the dashboard undertaken by advisers found that non-experts perceived LADA to be more useful compared to experts.

Hellings and Haelermans (2020) designed a dashboard aiming to provide study progress updates to students along with their predicted probability of success in a course accompanied with the predicted grade. A weekly email with the dashboard link was sent to the participants with the intention of encouraging them to use the dashboard frequently. Linear models were used to predict the grade mark, while AdaBoost predicted the course outcomes. Dashboard usage was correlated with a positive impact on student online engagement, but no impact on final exam grades and course completions was detected.

Plak et al. (2021) evaluated an EWS with student counsellors to estimate the effect of counselling on first-year student dropout rates and academic performance. To predict student dropout, the authors used SVM, RF, and LR models. LR outperformed all the other algorithms in this context. The findings indicated that EWS-assisted counselling did not reduce dropout or increase the credits obtained by the end of the academic year. It was hypothesized that this was due to non-targeted and non-actionable feedback/recommendations being provided to at-risk students.

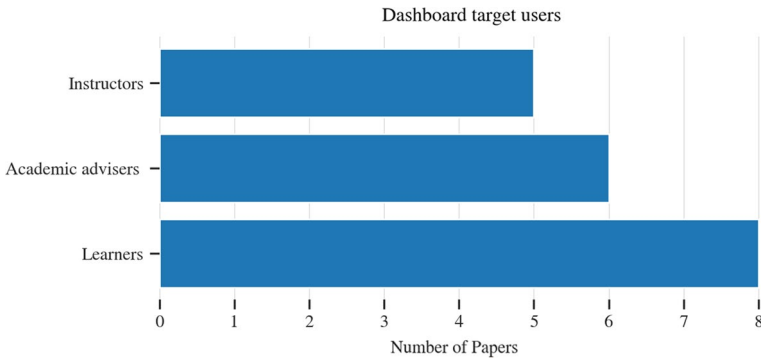


Fig. 3 Target users of the dashboards. Some dashboards target multiple audiences

6 Results

This section summarizes the key areas of interest that emerged from the 15 selected studies. We present them in the following order, key themes representing the overarching purpose of the LADs, target users, data types used, prediction methods, visualization types, evaluation techniques and finally study limitations.

6.1 Key Themes

The 15 dashboards have been categorized into two overarching themes. Eleven dashboards were used to track study planning. Their primary aim was to support academic advisers or instructors in assisting students during courses and in more effectively planning their subsequent course selections. Study planning advice was provided by conducting a session with students or sending them an awareness message with guidance during an ongoing course. With majority of studies focusing on either providing study advice or sending awareness messages to at-risk students, we found one study by Corrigan et al. (2015), that delivered advice messages to both at-risk and not-at-risk students.

The second theme centered on monitoring learning progress from student learning activities. Four studies aimed at improving learning outcomes by increasing student online engagement.

6.2 Target Users

Our study found that the LADs target learners, instructors and academic advisers, with learners being the most common target group (as shown in Fig. 3). In 11 cases, the LADs targeted a single user group. In the remaining four, two targeted both instructors and learners while the other two targeted instructors and academic advisers.

6.3 Types of Data Used for Prediction

Data for enabling the LADs were collected from a variety of sources including LMSs, demographics, pre-academic, library interaction as well as assessment grades, with studies using various combinations of these inputs.

LMSs provide online learning digital footprints like logs of interactions such as forum activities (messages viewed and posted), durations of sessions, quizzes taken, resources accessed etc. Most studies utilized logs of these interactions as absolute counts for making onward predictions, without normalizing them by representing values for each student in relation to those of their cohort. Pre-academic data comprised students' academic background information such as high school GPAs, past academic history, aptitude tests and enrolment details, while characteristics of the demographic data were age, ethnicity, and gender. In addition to these data, most of the researchers used assessment grades (quiz scores, in-between assignment grades, course grades) made available during the semester progression as well as prior semesters' grades.

6.4 Prediction Methods

Our review reveals a wide range of machine learning algorithms being utilized for predictive analytics (Table 1). These algorithms differ in various aspects, such as in their inherent biases and assumptions, their ability to generalize beyond training data and the amount of data required for training the models (Osmanbegovic & Suljic, 2012). Sweeping claims that one algorithm is superior to another across all datasets in this domain are unsupported.

Algorithms include DTs (Quinlan, 1986) which use partition rules to divide the data into groups based on a single variable. The process continues until all the variables have been utilized (Ferreira et al., 2001). The other tree-based classifiers commonly used for predictions are CART (Breiman et al., 1984) and RF (Breiman, 2001). CART has robust mechanisms for pruning large DTs in order to reduce overfitting which compromises generalization. While RF is based on combining multiple DTs into a single model also known as ensemble methods. Several studies have leveraged ensemble-based algorithms as they tend to produce classifiers having more robust generalizability mechanisms.

Logistic Regression (LR) (Cole, 1991) and SVM (Cortes et al., 1995) are examples of function-based algorithms that extract knowledge that is encoded in the form of mathematical functions, with LR being used in more studies compared to SVM, possibly due to LR being easier to specify in its tunable parameters. Two papers used Neural Networks (NNs) for predicting student risk levels.

LADA was the only study that used clustering which is an unsupervised learning technique. Clustering automatically groups individuals into distinctive groups based on commonalities. Once a cluster model is generated, new data points can subsequently be assigned to a cluster based on some distance criterion (Ochoa, 2016).

Ten studies utilized predictive analytics for forecasting the final grades with the aim of identifying at-risk students. One study relied on predictive analytics to determine whether students would submit their upcoming assignments. Another attempted to predict whether students would return to complete their studies in the following semester, while one study generated an estimation regarding the expected duration until qualification completion for individual students. One study sought to assist students in selecting course options by predicting their performance outcomes across different options.

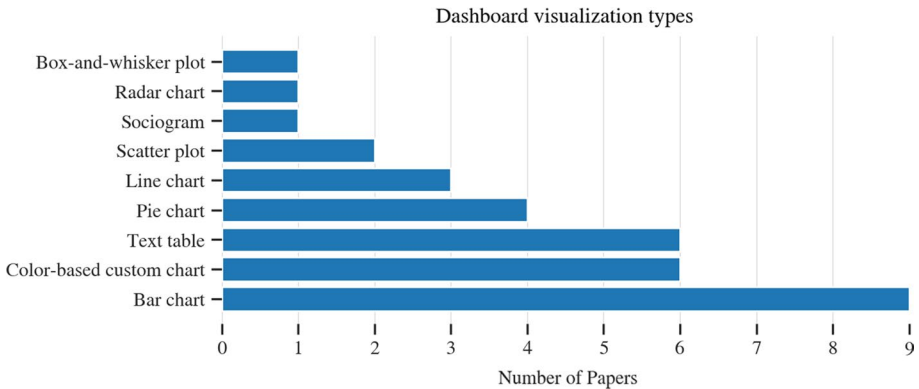


Fig. 4 Visualization types (We coalesced some chart types that rely on identical visual encoding channels into more general categories. Donut and Gauge charts have been placed into the Pie Chart category. Slider, Win-loss, Histograms and Stacked Bar Graphs have been grouped into a Bar Graph category. Color-based visualizations that which are non-standard and employ custom displays with an overarching emphasis on color as a visual encoding channel, have been grouped into Color-based Custom Charts. This category includes visualisations like traffic signals and risk-quadrant.)

However, while all the reviewed LADs utilized predictive analytics, none provided model interpretability to explain how the model works overall, nor were there any instances of insights being directly communicated about model reasoning (such as how a particular student was predicted to have a given outcome) making it difficult to recommend a specific remediation or even to trust the system.

6.5 Evaluating the Prediction Accuracy

Once predictive models are created, it is crucial to evaluate their generalizability on new (test) data. A variety of measures can be used for this task, and it is good practice to assess models using a suite of evaluation metrics. Common metrics in literature are accuracy, F-measure, mean absolute error (MAE), precision, recall, false positive (FP) and false negative (FN) rates, and receiver operating characteristic (ROC) curves. Accuracy is simply defined as a proportion of correct classifications. Evaluating a model with accuracy alone can be misleading if a dataset is imbalanced between samples belonging to different labels (outcomes). Majority of studies combined accuracy with other metrics in determining the efficacy of their models. However, two studies (Essa & Ayad, 2012; Wang et al., 2018) used accuracy alone. Given that the cost of misclassification is different for those who are at-risk as opposed to those who are not, it is also important to consider the model behavior on each of the two groups separately by using measures like recall, precision, and FP/FN rates. This was followed by some studies.

Few studies used the F-measure which is the harmonic mean of precision, and recall, and is arguably a more informative evaluation metric compared to accuracy for imbalanced datasets. MAE which is the difference between the model's predicted probability of non-completion and the mean of actual non-completion students was used by Plak et al. (2021). Meanwhile, the more advanced ROC curve that represents a classifier's performance on a set of performance thresholds balancing recall and FP rates (Fawcett, 2006) was used by three studies. A summary of these evaluation measures is shown in Table 1.

6.6 Visualization Methods

Graphical data representation is a key consideration in dashboards, since the main purpose of data visualization is effective communication. When designed correctly, data visualizations can provide support to students, teachers or educational administrators (Buenaño-Fernández et al., 2019). Figure 4 depicts the frequency of usage of various visualization charts. The most widely used visualizations are bar charts (Sahin & Ifenthaler, 2021). Bar charts encode quantitative data effectively by relying on the ability of the human visual cortex to accurately distinguish relative differences in lengths of graphical cues (Mackinlay, 1986). Other visualization techniques which are accurate at communicating quantitative information are line (capturing trends), scatter (conveying relative positions) and box-and-whisker plots (showing distributions), though their usage is not as widespread.

Color-based custom charts representing graphs which rely solely on color as the visual encoding channel to communicate information are also used extensively. Examples of these charts are traffic signals and risk-quadrants aiming at communicating nominal data. In all instances, these have been used to convey risk categories. Tables containing text-based information have also been widely used. Tables provide an effective and powerful way to present data meaningfully (Midway, 2020).

Various forms of pie charts, including donut and gauge plots were also used frequently. However, pie charts are problematic when the aim is to precisely convey quantitative information due to their reliance on humans being able to reliably quantify slopes and differences in angles. Angles and slopes are poor visual encoding channels and LADs utilizing them will have met their goals sub-optimally. Two studies used the gauge plot to present risk levels to students. While this is not an optimal visualization technique, given the simplicity of the information being communicated, arguably these charts were likely adequate. One instance shows the use of sociograms for presenting online communication and networking data as well as one use of a radar chart; the latter exhibiting problematic characteristics for interpretation due to its reliance on the reader being able to accurately quantify differences between areas of irregular shapes. Three studies did not mention details regarding their visualizations and are not included in Fig. 4.

6.7 Dashboard Evaluations

Our findings reveal that evaluation approaches to measure dashboard effectiveness were followed with variable rigor among the reviewed papers (see Table 1). Most of the studies did not mention any evaluation criteria nor conduct any evaluations. Of the remaining studies, the majority used mixed methods that combined qualitative and quantitative methods for evaluating the usefulness of the dashboards. One study conducted a qualitative approach by evaluating general aspects like LAD usability by inquiring into users' perceptions of their interaction with the tool.

Four studies investigated the effects of LADs on students' final outcomes. Of these, Arnold and Pistilli (2012), and Jayaprakash et al. (2014) found that providing risk-level predictions to students helped them achieve better grades in their courses. Bañeres et al. (2020) claimed that usage of the LAD showed some positive impact on students' final outcomes, but they could not determine whether this was attributable to the utilization of the tool or to interventions. Hellings and Haelermans (2020) claimed that LAD usage showed

positive effects on the student online engagement, although no similar effect was found for students' final exam performances.

7 Discussion

Our research shows that until now only a limited number of dashboards have incorporated predictive analytics. However, the trend is increasing (Fig. 2). The few which have implemented and evaluated this technology have mostly reported positive effects. A proportion of the positive effects can undoubtedly be ascribed to associated interventions initiated with at-risk students, raised by predictive models' early warnings. However, some effects can likely also be attributed to LAD-embedded predictive analytics and the self-reflection that they trigger in at-risk students, which result in positive behavioral adjustments. While it is not possible at this stage to disentangle the contribution that different factors have on the effects in eliciting positive outcomes, the utilization of predictive models for real-time identification of at-risk students remains a forward-thinking strategy that has so far indicated to hold promise as a tool for enhancing student success and retention rates.

Many of the reviewed articles were used to track study plans for at-risk learners but did not provide any direct insights into detailed causes behind the models' risk predictions, thus making it difficult to offer learners a tailored set of remedial actions. In the case of Course Signals, Tanes et al. (2011) performed content analysis of the feedback messages sent by instructors to students after receiving alerts. The authors noted the lack of instructive or process feedback types in the instructor messages sent to students, meaning that the students did not know what specific behavioral adjustments needed to take place. Therefore, alongside accurate predictions, it appears important to consider how prediction mechanics and model outputs are presented as advice to at-risk students.

Predictive analytics, when performed well, presents fresh opportunities in terms of the timing of intervention strategies. Machine learning can detect subtle yet complex multi-dimensional patterns early on in a semester that humans cannot, enabling an early initiation of potentially more effective interventions. However, determining an optimal time to intervene with identified at-risk students is somewhat of a challenge since predictive models improve in their accuracy with more data as a semester progresses and each students' digital footprint increases. Most of the studies have used black-box machine learning algorithms. Unfortunately, researchers in the LAD field have not yet drawn from emerging machine learning technologies that convert uninterpretable black-box models into ones that are understandable, and which consequently offer students and instructors deeper insights about what are the key drivers of negative predicted outcomes for a given student.

Jayaprakash et al. (2014) advise that predictive models do not influence course completion and retention rates unless they are combined with effective intervention strategies aimed at supporting at-risk students. This could again be due to the lack of interpretability of the predictive models and the absence of prediction explainability to learners as to how exactly the models arrived at given conclusions (Mathrani et al., 2021). Model interpretability and explainability can be realized with the use of the Shapely Additive Explanations (SHAP) method (Lundberg et al., 2021), Local Interpretable Model-agnostic Explanations (LIME (Ribeiro et al., 2016)) or anchors (Ribeiro et al., 2018) which have recently become popular.

With recent advances, current machine learning technologies are able to not only explain models and their predictions but are also able to offer data-driven and automated

counterfactuals which offer prescriptive capabilities that can clearly articulate to the learners what behavioral adjustments would hypothetically result in future positive predictive outcomes (Susnjak et al., 2022). As it currently stands, existing LADs only display the predicted outcomes, indicating significant gaps and pointing towards rich future research opportunities.

Findings from the reviewed papers revealed a considerable diversity in the utilization of predictive analytics. Unsurprisingly, majority of reviewed articles integrated predictive analytics for forecasting the final course academic performance. Others saw the state of being at-risk as being broader, and used proxies such as assessment submission prediction, likelihood of returning to complete current studies and the prediction of the overall qualification completion as alternate approaches to formulate the problem.

LADs frequently aim to communicate quantitative information which requires precision and, as alluded to earlier, an appropriate matching between the data types and the visual encoding channels (Munzner, 2008). Unsuitable visualization design choices add to the cognitive load of users and fail in their original intent, and equally this occurs when color is used indiscriminately in dashboards (Bera, 2016). Visualization theory has firmly established the primacy of various visualization encoding channels (e.g., position, length, slope/angle, area, depth, hue/tint, shape, curvature, volume etc.) when communicating data, whether it be for continuous, ordinal, or categorical types (Munzner, 2014). We observed in our survey that a number of studies used graphing components like donut graphs, gauge plots, pie and radar charts. These rely on visual encoding channels like slope/angle and area which are suboptimal for communicating information, be it quantitative or qualitative. This raises questions about the degree to which visualization theory is being drawn upon in the design of some recent LADs. The issue of inappropriate visualizations specifically within LADs has already been raised. Park and Jo (2019) add that unclear visualizations can restrict the student's sensemaking of the LADs. While Schwendimann et al. (2017) posit that the granularity of information being displayed on the dashboard should be determined with proper visualization techniques so that users are not confused or overwhelmed by the amount of information presented to them. Our study shows that 40% of the dashboards use three or more different visualization techniques, with one dashboard using seven which likely leads to cognitive overload rather than clarity. Midway (2020) also mentioned that effective visualizations foster in the audience an intended understanding and interpretation of the data, while the reverse holds.

It is also the case that different users have different informational needs, and the way they perceive messages influences their actions. Hence, it is recommended that the user perspective is kept at the forefront during dashboard design. Perception of learning varies among users (Verbert et al., 2013a, 2013b); therefore, LAD designers must appreciate how individuals interact with visuals, and what impact visuals have on their learning processes (Bodily & Verbert, 2017; Corrin & de Barba, 2014).

The evaluation of the impact of LADs towards learner outcomes is an area that requires much attention. Studies which had conducted evaluations were mainly focused on the functionality and usability of LADs, and while they highlighted the potential effect on learning, they did not demonstrate quantifiable effects of LADs as a pedagogical treatment. Though conducting usability testing has merit, future studies should also assess whether LAD usage has positive effects on students' final outcomes. This can be determined by conducting an experimental study among two groups (control and treatment) over some length of time in order to determine if there is a difference in outcomes. We also found that most of the LADs were working prototypes; hence, students were unable to use them in real course settings and provide reliable feedback on their effectiveness.

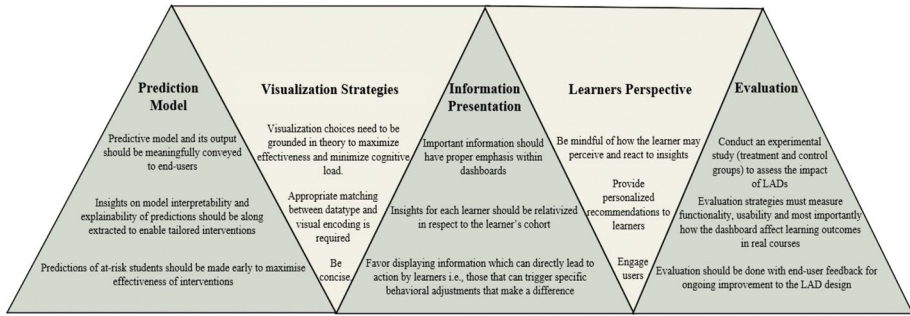


Fig. 5 Recommendations for Learning Analytics Dashboard Design

8 Recommendations for LAD Design

The key points that have emerged from our review of the 15 selected LADs are the necessity of fully leveraging machine learning for predictive and prescriptive analytics, the need to design dashboard displays with information visualization theory in mind, having a focus on appropriately presenting information to end-users, adopting a learner's perspective and conducting an evaluation strategy (refer Fig. 5). Presenting mere outputs of prediction models on their own within LADs are unlikely to result in significantly improved learning outcomes. The information accompanying students' predicted outcomes needs to include model explainability and ideally, prescriptive analytics which can support a tailored intervention strategy to assist at-risk students. Next, visualization choices need to be grounded in established theory and best practice in order to maximize effectiveness and minimize cognitive load. Further, being mindful of how learners perceive information is key to the effectiveness of the message that is being conveyed. Therefore, strategies to engage users (via personalized recommendations and other targeted motivating messages) should be at the forefront. Finally, institutions should have some evaluation strategy to ensure that the LADs meet the goals for which they have been designed. We propose that such evaluations should be done with end-users (i.e., learners, instructors and academic advisers) for ongoing improvements to the LAD designs. Figure 5 frames our recommendations for LAD design, deployment and its effective utilization as a technological tool possessing a pedagogical value.

9 Conclusion

This article has conducted an extensive literature review on various LADs that tracked student data and reported insights to various stakeholders. LADs leveraging predictive analytics were considered in this study. Our review indicates significant limitations in the current positioning of LADs in real-world educational settings. Our analysis finds that evidence of LADs' effectiveness to impact student learning outcomes is inconclusive. Current LAD evaluations are mostly limited to functionality and usability aspects only.

The reviewed LADs predominately aim at facilitating study planning and monitoring student learning progress, while targeting learners, academic advisers and instructors. Given the importance of effective visual presentation, our study finds that there is some

improvement to be made in terms of correct usage of visualization techniques in order to enhance precision and to reduce cognitive load.

We also find that predictive analytics are only now beginning to emerge as a tool within LADs. Our study concludes that its usage lacks maturity and lags in leveraging state-of-the-art machine learning technologies which enable model interpretability and explainability. We maintain that predictions of student's outcomes should be accompanied with outputs that provide high-level information about the mechanics of the models, as well as explanations of how a specific prediction was derived for a given student. Ideally, we should now also be starting to see the emergence of prescriptive analytics within LADs that utilize counterfactuals which can provide at-risk learners with actionable insights.

Our survey has enabled us to identify strengths and gaps in current LAD studies, and to formulate a framework comprising a set of recommendations for future LAD designs. We believe that adhering to the proposed LAD framework can aid educational providers in implementing more effective LADs, which facilitate improved learning outcomes.

Funding Open Access funding enabled and organized by CAUL and its Member Institutions. The authors have not disclosed any funding.

Data Availability The data is available to qualified researchers by request to the corresponding author.

Declarations

Conflict of interest The authors have not disclosed any competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Afzaal, M., Nouri, J., Zia, A., Papapetrou, P., Fors, U., Wu, Y., Li, X., & Weegar, R. (2021). Explainable AI for data-driven feedback and intelligent action recommendations to support students self-regulation. *Frontiers in Artificial Intelligence*, 4, 1–20. <https://doi.org/10.3389/frai.2021.723447>
- Agnihotri, L., & Ott, A. (2014). Building a student at-risk model: An end-to-end perspective. In: *Proceedings of the 7th international conference on educational data mining* (pp. 209–212).
- Arnold, K. E., & Pistilli, M. D. (2012). Course signals at Purdue. In: *Proceedings of the 2nd international conference on learning analytics and knowledge—LAK '12*. <https://doi.org/10.1145/2330601.2330666>
- Baker, R. S. (2016). Stupid tutoring systems, intelligent humans. *International Journal of Artificial Intelligence in Education*, 26(2), 600–614. <https://doi.org/10.1007/s40593-016-0105-0>
- Bañeres, D., Rodríguez, M. E., Guerrero-Roldán, A. E., & Karadeniz, A. (2020). An early warning system to detect at-risk students in online higher education. *Applied Sciences*, 10(13), 1–28. <https://doi.org/10.3390/app10134427>
- Bera, P. (2016). How colors in business dashboards affect users' decision making. *Communications of the ACM*, 59(4), 50–57. <https://doi.org/10.1145/2818993>

- Bodily, R., & Verbert, K. (2017). Review of research on student-facing learning analytics dashboards and educational recommender systems. *IEEE Transactions on Learning Technologies*, *10*(4), 405–418. <https://doi.org/10.1109/TLT.2017.2740172>
- Breiman, L. (2001). Random forests. *Machine Learning*, *45*(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- Breiman, L., Friedman, J. H., Olshen, R. A., & Stone, C. J. (1984). *Classification and Regression Trees*. Wadsworth, Inc.
- Buenaño-Fernández, D., Gil, D., & Luján-Mora, S. (2019). Application of machine learning in predicting performance for computer engineering students: A case study. *Sustainability*, *11*(10), 1–18. <https://doi.org/10.3390/su11102833>
- Charleer, S., vande Moere, A., Klerkx, J., Verbert, K., & de Laet, T. (2018). Learning analytics dashboards to support adviser-student dialogue. *IEEE Transactions on Learning Technologies*, *11*(3), 389–399. <https://doi.org/10.1109/TLT.2017.2720670>
- Cole, T. J. (1991). Applied logistic regression. *Statistics in Medicine*, *10*(7), 1162–1163. <https://doi.org/10.1002/sim.4780100718>
- Corrigan, O., Smeaton, A. F., Glynn, M., & Smyth, S. (2015). Using educational analytics to improve test performance (pp. 42–55). https://doi.org/10.1007/978-3-319-24258-3_4
- Corrin, L., & de Barba, P. (2014). Exploring students' interpretation of feedback delivered through learning analytics dashboards. In: *Proceedings of the Ascilite 2014 conference* (pp. 629–633).
- Cortes, C., Vapnik, V., & Saitta, L. (1995). Support-vector networks editor. In *Machine learning* (Vol. 20, pp. 273–297). Kluwer Academic Publishers.
- Delen, D., & Demirkan, H. (2013). Data, information and analytics as services. *Decision Support Systems*, *55*(1), 359–363. <https://doi.org/10.1016/j.dss.2012.05.044>
- Essa, A., & Ayad, H. (2012). Student success system: risk analytics and data visualization using ensembles of predictive models. In: *Proceedings of the 2nd international conference on learning analytics and knowledge—LAK '12* (pp 158–161). <https://doi.org/10.1145/2330601.2330641>
- Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recognition Letters*, *27*(8), 861–874. <https://doi.org/10.1016/j.patrec.2005.10.010>
- Ferreira JTA, Denison DG, Hand DJ (2001) Data mining with products of trees. In: *International symposium on intelligent data analysis* (pp 167–176). Berlin: Springer
- Greller, W., & Drachsler, H. (2012). Translating learning into numbers: A generic framework for learning analytics. *Journal of Educational Technology & Society*, *15*(3), 42–57.
- Gutiérrez, F., Seipp, K., Ochoa, X., Chiluíza, K., de Laet, T., & Verbert, K. (2020). LADA: A learning analytics dashboard for academic advising. *Computers in Human Behavior*, *107*(105826), 1–13. <https://doi.org/10.1016/j.chb.2018.12.004>
- He, L., Levine, R. A., Bohonak, A. J., Fan, J., & Stronach, J. (2018). Predictive analytics machinery for STEM student success studies. *Applied Artificial Intelligence*, *32*(4), 361–387. <https://doi.org/10.1080/08839514.2018.1483121>
- Hellings, J., & Haelermans, C. (2020). The effect of providing learning analytics on student behaviour and performance in programming: A randomised controlled experiment. *Higher Education*. <https://doi.org/10.1007/s10734-020-00560-z>
- Hu, Y.-H., Lo, C.-L., & Shih, S.-P. (2014). Developing early warning systems to predict students' online learning performance. *Computers in Human Behavior*, *36*, 469–478. <https://doi.org/10.1016/j.chb.2014.04.002>
- Jayaprakash, S. M., Moody, E. W., Lauría, E. J. M., Regan, J. R., & Baron, J. D. (2014). Early alert of academically at-risk students: An open source analytics initiative. *Journal of Learning Analytics*, *1*, 1–42. <https://doi.org/10.18608/jla.2014.11.3>
- Kitchenham, B., & Charters, S. (2007). Guidelines for performing systematic literature reviews in software engineering, pp. 1–65.
- Kuzilek, J., Hlosta, M., Herrmannova, D., Zdrahal, Z., & Wolff, A. (2015). OU Analyse: analysing at-risk students at The Open University. In: *Learning analytics and knowledge conference* (pp. 1–16).
- Lundberg, S. M., Allen, P. G., Lee, S.-I. (2021). A unified approach to interpreting model predictions. Available online: <https://github.com/slundberg/shap>. Accessed 12 Oct 2021
- Mackinlay, J. (1986). Automating the design of graphical presentations of relational information. *ACM Transactions on Graphics*, *5*(2), 110–141. <https://doi.org/10.1145/22949.22950>
- Matcha, W., Uzir, N. A., Gasevic, D., & Pardo, A. (2020). A systematic review of empirical studies on learning analytics dashboards: A self-regulated learning perspective. *IEEE Transactions on Learning Technologies*, *13*(2), 226–245. <https://doi.org/10.1109/TLT.2019.2916802>

- Mathrani, A., Susnjak, T., Ramaswami, G., & Barczak, A. (2021). Perspectives on the challenges of generalizability, transparency and ethics in predictive learning analytics. *Computers and Education Open*, 2(100060), 1–9. <https://doi.org/10.1016/j.caeo.2021.100060>
- Midway, S. R. (2020). Principles of effective data visualization. *Patterns*, 1(9), 100141. <https://doi.org/10.1016/j.patter.2020.100141>
- Munzner, T. (2008). LNCS 4950—Process and pitfalls in writing information visualization research papers. <http://www.cs.ubc.ca/~tmm>
- Munzner, T. (2014). *Visualization analysis and design*. CRC Press.
- Ochoa, X. (2016). Adaptive multilevel clustering model for the prediction of academic risk. In: *XI Latin American conference on learning objects and technology (LACLO), 2016* (pp. 1–8). <https://doi.org/10.1109/LACLO.2016.7751800>
- Osmanbegovic, E., & Suljic, M. (2012). Data mining approach for predicting student performance. *Economic Review: Journal of Economics and Business*, 10(1). <http://hdl.handle.net/10419/193806>
- Park, Y., & Jo, I.-H. (2019). Factors that affect the success of learning analytics dashboards. *Educational Technology Research and Development*, 67(6), 1547–1571. <https://doi.org/10.1007/s11423-019-09693-0>
- Plak, S., Cornelisz, I., Meeter, M., & Klaveren, C. (2021). Early warning systems for more effective student counselling in higher education: Evidence from a Dutch field experiment. *Higher Education Quarterly*. <https://doi.org/10.1111/hequ.12298>
- Podgorelec, V., & Kuhar, S. (2021). Taking advantage of education data: Advanced data analysis and reporting in virtual learning environments. *Electronics and Electrical Engineering*, 114(8), 111–116. <https://doi.org/10.5755/j01.eee.114.8.708>
- de Quincey, E., Briggs, C., Kyriacou, T., & Waller, R. (2019). Student centred design of a learning analytics system. In: *Proceedings of the 9th international conference on learning analytics & knowledge* (pp 353–362). <https://doi.org/10.1145/3303772.3303793>
- Quinlan, J. R. (1986). Induction of decision trees. *Machine Learning*. <https://doi.org/10.1007/BF00116251>
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). Why Should I Trust You? In: *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 1135–1144). <https://doi.org/10.1145/2939672.2939778>
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2018). Anchors: High-precision model-agnostic explanations. In: *The thirty-second AAAI conference on artificial intelligence* (pp. 1527–1535). www.aaai.org. <https://ojs.aaai.org/index.php/AAAI/article/view/11491>
- Sahin, M., & Ifenthaler, D. (2021). Visualizations and dashboards for learning analytics: A systematic literature review. *Visualizations and Dashboards for Learning Analytics*. https://doi.org/10.1007/978-3-030-81222-5_1
- Schwendimann, B. A., Rodriguez-Triana, M. J., Vozniuk, A., Prieto, L. P., Boroujeni, M. S., Holzer, A., Gillet, D., & Dillenbourg, P. (2017). Perceiving learning at a glance: A systematic literature review of learning dashboard research. *IEEE Transactions on Learning Technologies*, 10(1), 30–41. <https://doi.org/10.1109/TLT.2016.2599522>
- Susnjak, T., Ramaswami, G. S., & Mathrani, A. (2022). Learning analytics dashboard: A tool for providing actionable insights to learners. *International Journal of Educational Technology in Higher Education*, 19(1), 1–23. <https://doi.org/10.1186/s41239-021-00313-7>
- Tanes, Z., Arnold, K. E., King, A. S., & Remnet, M. A. (2011). Using signals for appropriate feedback: Perceptions and practices. *Computers & Education*, 57(4), 2414–2422. <https://doi.org/10.1016/j.compedu.2011.05.016>
- Umer, R., Susnjak, T., Mathrani, A., & Suriadi, L. (2021). Current stance on predictive analytics in higher education: Opportunities, challenges and future directions. *Interactive Learning Environments*. <https://doi.org/10.1080/10494820.2021.1933542>
- Valle, N., Antonenko, P., Dawson, K., & Huggins-Manley, A. C. (2021). Staying on target: A systematic literature review on learner-facing learning analytics dashboards. *British Journal of Educational Technology*, 52(4), 1724–1748. <https://doi.org/10.1111/bjjet.13089>
- Verbert, K., Duval, E., Klerkx, J., Govaerts, S., & Santos, J. L. (2013a). Learning analytics dashboard applications. *American Behavioral Scientist*, 57(10), 1500–1509. <https://doi.org/10.1177/0002764213479363>
- Verbert, K., Govaerts, S., Duval, E., Santos, J. L., van Assche, F., Parra, G., & Klerkx, J. (2013b). Learning dashboards: An overview and future research opportunities. *Personal and Ubiquitous Computing*. <https://doi.org/10.1007/s00779-013-0751-2>
- Wang, Z., Zhu, C., Ying, Z., Zhang, Y., Wang, B., Jin, X., & Yang, H. (2018). Design and implementation of early warning system based on educational big data. In: *2018 5th international conference on systems and informatics (ICSAI)*, pp. 549–553. <https://doi.org/10.1109/ICSAI.2018.8599357>

Yoo, Y., Lee, H., Jo, I.-H., & Park, Y. (2015). Educational dashboards for smart learning: Review of case studies. In: *International conference on smart learning environments* (pp. 145–155). Springer. https://doi.org/10.1007/978-3-662-44188-6_21

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Chapter 3

Predicting Students Final Academic performance using Feature Selection Approaches

The contents of this chapter are from the following article. In accordance with the IEEE republishing policy, any full text that has been included, is the unmodified accepted article.

@2020 IEEE. Reprinted, with permission, from: Ramaswami, G. S., Susnjak, T., Mathrani, A., & Umer, R. (2020, December). Predicting Students Final Academic Performance using Feature Selection Approaches. In *2020 IEEE Asia-Pacific Conference on Computer Science and Data Engineering (CSDE)* (pp. 1-5). IEEE. doi:10.1109/CSDE50874.2020.9411605.

The IEEE open access policy permits the use, sharing, adaptation, distribution, and reproduction in any medium or format as long as appropriate credit goes to the authors. IEEE does not endorse any of Massey University's products or services. Internal or personal use of this material is permitted. If interested in reprinting/republishing IEEE copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribution, please go to: http://www.ieee.org/publications_standards/publications/rights/rights_link.html to learn how to obtain a license from the correct link.

STATEMENT OF CONTRIBUTION DOCTORATE WITH PUBLICATIONS/MANUSCRIPTS

We, the candidate and the candidate's Primary Supervisor, certify that all co-authors have consented to their work being included in the thesis and they have accepted the candidate's contribution as indicated below in the *Statement of Originality*.

Name of candidate:	Gomathy Suganya Ramaswami
Name/title of Primary Supervisor:	Dr. Teo Susnjak
In which chapter is the manuscript /published work:	Chapter 3
Please select one of the following three options:	
<input checked="" type="radio"/> The manuscript/published work is published or in press <ul style="list-style-type: none"> • Please provide the full reference of the Research Output: Ramaswami, G. S., Susnjak, T., Mathrani, A., & Umer, R. (2020, December). Predicting Students Final Academic Performance using Feature Selection Approaches. In 2020 IEEE Asia-Pacific Conference on Computer Science and Data Engineering (CSDE) (pp. 1-5). IEEE. 10.1109/CSDE50874.2020.9411605. 	
<input type="radio"/> The manuscript is currently under review for publication – please indicate: <ul style="list-style-type: none"> • The name of the journal: • The percentage of the manuscript/published work that was contributed by the candidate: • Describe the contribution that the candidate has made to the manuscript/published work: 	
<input type="radio"/> It is intended that the manuscript will be published, but it has not yet been submitted to a journal	
Candidate's Signature:	Gomathy Ramaswami <small>Digitally signed by Gomathy Ramaswami Date: 2022.11.07 07:42:59 +13'00'</small>
Date:	07-Nov-2022
Primary Supervisor's Signature:	Teo Susnjak <small>Digitally signed by Teo Susnjak Date: 2022.11.07 15:30:10 +13'00'</small>
Date:	7-Nov-2022

This form should appear at the end of each thesis chapter/section/appendix submitted as a manuscript/ publication or collected as an appendix at the end of the thesis.

Predicting Students Final Academic Performance using Feature Selection Approaches

Gomathy Suganya Ramaswami
*School of Natural and Computational
Sciences*
Massey University
Auckland, New Zealand
G.Ramaswami@massey.ac.nz

Teo Susnjak
*School of Natural and Computational
Sciences*
Massey University
Auckland, New Zealand
T.Susnjak@massey.ac.nz

Anuradha Mathrani
*School of Natural and Computational
Sciences*
Massey University
Auckland, New Zealand
A.S.Mathrani@massey.ac.nz

Rahila Umer
Engineering and Management Science
*Balochistan University of Information
Technology*
Quetta, Pakistan
rahilaumer@buitms.edu.pk

Abstract— Data analytics and data mining techniques are increasingly being used to assist managers, business leaders, and policymakers in their decision-making. These techniques help them distill the useful information hidden in their ever-growing databases that are distributed across various data stores. Another area in which these techniques are utilized is the education domain, wherein data analytics can be leveraged to predict students' academic performance. Teachers will therefore be better informed in providing their students with additional support if required. Learning Management Systems (LMS) captures student online learning data and are extensively used by researchers to gather data-driven insights on students' overall academic performance. This study aims to predict students' academic performance with the LMS data from an online training course using various machine learning algorithms. The study further highlights that selection of features by using feature selection methods will enable early prediction of student academic performance. Early predictions can benefit those students who are considered to be at risk of failing a course with targeted early interventions to help improve their performance throughout the course.

Keywords—learning dashboard, feedback, student feedback system, visualization tool, dashboard

I. INTRODUCTION

Web-based education or e-learning has grown exponentially over the last few years. In particular, learning management systems (LMS) that emerged from the field of e-learning, have become an integral part of instructional delivery practice, with universities supplementing face-face learning with LMS [1]. LMS have enabled students to personalize their learning, review specific course material or engage in peer discussions. In doing so, students leave traces of their online interactions, and as such massive amounts of data are stored in the LMS. This data can be leveraged to provide rich data-driven insights on students' online behaviors since all their interactions within the LMS are captured.

Moodle is a widely used open-source learning management system. Moodle can accumulate vast amounts of user information that can be further used to conduct deeper analyses on student behaviors. To effectively use such data for improved decision-making, better approaches for extracting knowledge from large repositories are

required. Educational Data Mining (EDM) is an emerging discipline that can help researchers discover useful and distinct patterns within the voluminous data and create data-driven strategies to enhance the institutional teaching practices [2]. There is a wide range of EDM tasks, one of the most important being student performance prediction [3]. Predicting student performance enables instructors to take action when required, for example, weak students may be identified and provided with additional support. The main aim is to achieve higher levels of quality in education delivery and develop more personalized forms to impart education, even in large group settings [4].

The focus of this paper is to predict the final academic performance of students enrolled in a course using various machine learning classifiers. We are also interested in finding out the earliest timeframe in identifying the at-risk students of failing in a course. This study was conducted for a single course by collecting the activity log of the student interactions from the Moodle platform. We plan to subsequently extend the results of this study by including multiple courses in future analysis. Having posed the focus of this study, the rest of the paper is structured as follows. Section 2 discusses the various techniques and the data used in previous studies, followed in section 3 by a description of the data selection and machine learning methods adopted in this study. Section 4 explores the methodology adopted in this study, while sections 5 and 6 includes a discussion of the results and conclusions that can be drawn.

II. EXISTING REVIEWS

A review was conducted on recent studies, that is, those that were published between 2015 and 2020 and which had used LMS as their source of data for analysis. A study by Tempelaar, Rientities, and Giesbers [5] tested the predictive power of learning dispositions. They used demographic data, entry test data, LMS data, and formative assessment results of about 922 students, wherein Blackboard was used as the LMS tool. Their findings revealed that the at-risk students could be best detected using formative assessments and that learning performance predictions using the LMS data were not very accurate.

A similar study was conducted by Gasevic et al. [6] on a sample of LMS and with demographic data from 4134

students from nine undergraduate courses. The results showed that the inclusion of trace data from LMS to the student demographic data improved the accuracy of predicting students' marks from 5% to 16%. A multiple linear regression model was used to perform these tasks. The authors also found a significant difference across the courses in the association between the demographic data and trace data, ranging from 2% to 70% variability depending on the course.

Another study led by Amrieh, Hamtini and Alijarah [7] collected the LMS data from Kalboard 360 along with the demographic data and the academic background features. The authors applied the filter-method using an information gain-based selection algorithm for selecting features that are most important to build students' performance models. They extracted records from 480 students and used Artificial Neural Network (ANN), Naïve Bayesian (NB), and Decision tree (DT) for evaluating the student's performance and found ANN outperforms with 79.1% accuracy over the other classifiers. The obtained predictive model was tested using unlabeled newcomer students, and the achieved accuracy was more than 80%.

Konda et al. [8] were interested in building an automatic detection method to pinpoint the academically at-risk students using the LMS data. The data of 202 students were used for this study and well-known classifiers like logistic regression, random forest, and support vector machine was used to predict the GPA of the students. Of those tested, the random forest classifier was the most promising regarding stability, precision, and recall. When applied to the LMS data, their model was able to detect 40% of the at-risk students after three weeks of the course.

Quinn and Gray [9] investigated whether data derived from LMS can be used to predict the academic performance of the students in a blended learning environment. The data were extracted from 607 students for 29 classes. Random Forest, Gradient Boosting, k- Nearest Neighbors, and Linear Discriminant Analysis algorithms were evaluated, and the random forest was found to outperform all the other classifiers with 92% accuracy. The authors were aiming for early prediction and found that classifiers trained on the week 10 dataset did significantly better than classifiers trained on the week 6 dataset. The authors also suggested that the prediction accuracy might improve by including formative assessments along with the LMS data.

III. DATASETS

We conducted an exploratory research using student data that was extracted from courses offered at one Australasian university. These courses were delivered in a blended learning environment, that is they included a combination of e-learning and face-to-face learning [10]. Moodle has a built-in feature that can track student activity [11]. Therefore, activity logs from Moodle, assignment grades, and demographic data have been used for this study.

Table 1 describes the course information used in this study. The reason for choosing this particular course was due to its large student enrolments. Based on their predicted score for their final exam, students were classified into two groups: high-risk and low-risk. The gender of the student was coded to a numerical value of 0 for 'Male', and 1 for 'Female'. There were two assessments for this course.

Demographic data can be used as the only obtainable source of information for a new student when neither their previous results nor their activities are available in the LMS [12].

TABLE1: INFORMATION ABOUT THE COURSE

Course Name	Class Size		Number of Assessments	Student Grade Distribution	
	Male	Female		High-risk	Low-risk
Introduction to Finance	39	73	2	52	60

A. Moodle Action Logs

Fig 1 shows the various activities offered in the Moodle platform. The Moodle logs recorded information about each action that was performed by each individual student. Each time a student viewed a specific activity or participated in a particular activity, the information about that student and their activity was captured (refer Table 2). The most common activities that were logged pertained to quizzes, forums, assignments, and folder searches. Student usage of specific activities varied depending on the nature of each course. Since we were interested only in students' activities, therefore only the student data were extracted from log files for creating our dataset and activities logged by instructors which is not required were removed.




Fig. 1. Frequency of Events (activities) for a course

B. Ethics approval

Before conducting any study, it is important that researchers are aware of the ethical considerations and their obligations. This study has met the ethical requirements and has been approved by the Human Research Ethics Committee at the host university.

C. Tools Used

The study was conducted using Python and the relevant data processing and machine learning modules. Python has emerged as a popular tool for scientific computing tasks, which include the analysis and visualization of large datasets that are generally utilized for machine learning applications [13].

TABLE 2 ACTIONS PERFORMED BY A STUDENT IN MOODLE

Log Action	Description
URL viewed	Number of URLs viewed by the student
Book viewed	Number of books viewed by the student that are assigned to them on Moodle
Resource viewed	Total view of the source documents for a course that is uploaded on Moodle
Folders searched	Folders of the course searched by the student
Quiz viewed	Number of quizzes viewed by the student
Assignments viewed	Number of assignments viewed by the student
Assignments uploaded	Number of assignments uploaded
Assignments submitted	Total number of assignments submitted by the student
Quiz submitted	Number of quizzes submitted by the student
Quiz reviewed	Number of quizzes reviewed by the student
Forum created	Number of forum posts created by the student
Forum viewed	Number of forum posts viewed by the student
Forum searched	Number of forum posts searched by the student

IV. METHODOLOGY

The procedures for predicting the students’ academic performance are shown in the Fig 2. Data in raw form are generally not suitable for investigation, and particularly not for predictive analysis. Hence, pre-processing of the data was done first; this involved cleaning the data, filtering the non-related data, and finally organizing the data to obtain the finest mineable frame. Table 2 presents the set of log action types used in this study for analyzing the students’ online behavior.

The classification technique we used for prediction is one of the most widely used techniques in EDM. A training dataset in which all attributes are known is used to build the model. Then, a test dataset with unknown class attributes is used to verify the model. To classify the students into ‘High risk’ or ‘Low risk’, the following machine learning algorithms will be used: Random Forest (RF), Naïve Bayes (NB), Logistic Regression (LR) and k-Nearest Neighbors (kNN). These classifiers have shown good performance in previous studies too [14] [15] [16]. The performance of the classifiers was compared next using the hold-out method where the order of occurrences of the training–test split is maintained.

A. Model Evaluation Metrics

The model evaluation metrics were used to evaluate the goodness of fit between the model and the data [17]. F1-scores was used to evaluate the performance of the classifiers and it is commonly used in binary classification problems. F1-score is the harmonic mean of precision and recall.

$$F1\text{-score} = 2 * \frac{(\text{Precision} * \text{Recall})}{(\text{Precision} + \text{Recall})} \tag{1}$$

Precision belongs to the positive class and measures the positive prediction class (TP/(TP+FP)). Recall measures the number of correct predictions made from all positive predictions (TP/(TP+FN)) where TP denotes TruePositive, FP signifies FalsePositives and FN represents False Negatives.

B. Feature selection

It is not uncommon for the machine learning (ML) research to use all available features in a dataset without calculating the predictive value of each feature. Such a method might lead to overfitting and poor generalization [18]. However, identifying the ideal or near-optimal feature subset often results in better accuracies in ML. This can be performed using feature extraction, which extracts features from the raw dataset to train a model for machine learning algorithms [19] and is a fundamental task in data-preprocessing [7].

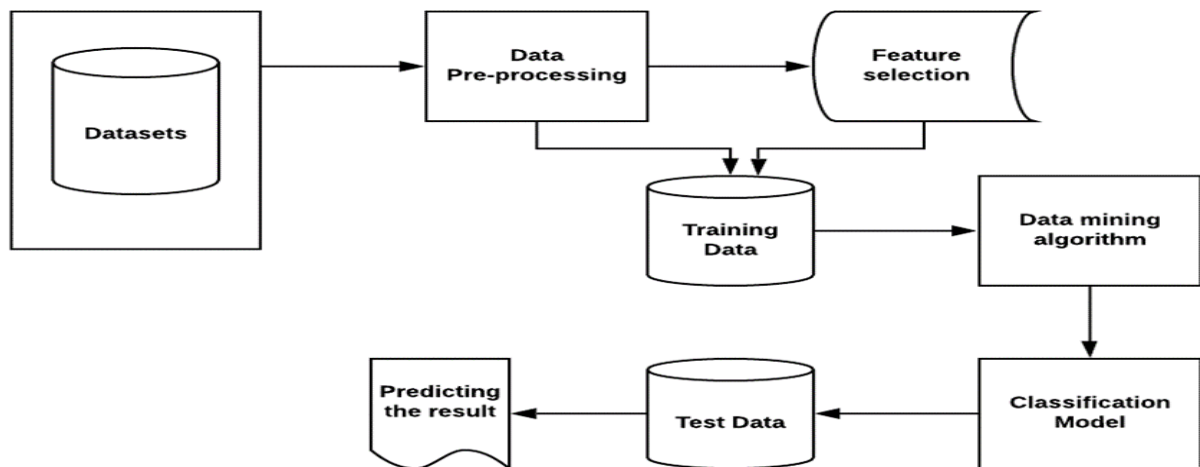


Fig. 2. Steps involved in prediction

Moreover, feature selection helps to reduce dimensionality, and realizes the association among features and feature values well [20]. The embedded method has been carried out for selecting the ideal features for prediction where the feature selection algorithm is integrated as part of the learning algorithm [21]. The table below (Table 3) shows the features selected using the feature selection method for prediction.

Experiment 1: Predicting student performance without the feature selection method.

Experiment 2: Predicting student performance using the feature selection method.

TABLE 3 FEATURES SELECTED USING EMBEDDED METHOD

Feature Name	Description	Importance
Assignment	The assignment score of the student	.1
Weekly quizzes	Weekly quizzes score of the student	.08
Assignment viewed	The number of times the student has viewed the assignment	.04
Quiz viewed	The number of times the student has viewed the quiz	.03
Age	Age of the student	.03
Folder viewed	The folder viewed by the student	.026
Book viewed	The number of books viewed by the student	.01

V. RESULTS AND DISCUSSION

As mentioned previously, the hold-out method was used, hence, the entire dataset was separated into training and testing. Seventy percent of the dataset was used for the training dataset to create the model to predict student performance, and the remaining 30% was used as the testing dataset to test the classification performance. The Logistic Regression model achieved the maximum F1-score compared to the other classifiers (Table 4). Hence, it was used for the early prediction to detect the at-risk students as soon as possible into the course.

TABLE 4 PERFORMANCE OF VARIOUS CLASSIFIERS

Classifiers	F1-score without FS	F1-score with FS
Naive Bayes	69.7%	79.41%
Random Forest	72%	79.41%
Logistic Regression	76.47%	85.29%
k-Nearest Neighbors	71.2%	78.6%

A. Experiment 1

To achieve this, the entire dataset was divided into weeks, and the data from week1 to week 16 were used for the prediction analysis. Data was used in a cumulative fashion for making predictions. For example, when making predictions for week 2 the students' data from week 1 were taken into consideration as well. This process was used until week 16. The two-assessment scores which happened on week 5 and week 10 respectively were also included along with the other data for prediction. The results of the experiment are shown in Fig 3. By looking into the graph, it can be seen that there is an improvement in F1-score over the weeks. But the F1-score was not continuously improving over the weeks and there is no consistency in F1-score.

B. Experiment 2

The same procedure as mentioned above were followed, but the features with high prediction accuracy were selected for the analysis to find out if there is any improvement in prediction accuracy. The graph in Fig 3 shows an improvement in accuracy compared to the previous method. The maximum F1-score of above 83% were achieved on week 11 and there was a steep growth in F1-score compared to the Experiment 1. As expected, selecting the features using the feature selection method helps in predicting the student's final performance better compared to the experiment 1.




Fig. 2. Logistic Regression performance with and without feature selection method

VI. CONCLUSION AND FUTURE SCOPE

Students interaction with an LMS results in generation of large datasets which are valuable for conducting this type of research. Many data mining techniques have proposed ways to extract hidden knowledge from these datasets. The extracted knowledge helps to improve the learning process of the students which leads to enhancing the performance of the students and the overall educational output. The current study focuses on predicting students' academic performance using a classification method where four different classifiers, namely RF, NB, LR, and k-NN. LR outperformed the other classifiers; hence, it was used for early prediction of students' final performance.

Two experiments were conducted, one in which the features were selected to give a high prediction accuracy and the other in which all the features were selected. The results showed that selecting the features using the feature selection method has better discriminative power. This illustrates feature selection is important to consider for predictive modelling which is however is not followed in most of the research study.

Future plans include evaluating the feasibility of the developed model using live data and testing the model on a larger sample size. Moreover, this study was limited in scope to developing the model using data drawn from only one course in one discipline. As an extension of the study we are interested in developing a more generalized model where multiple courses will be used to develop the model which will then be tested with data from other similar courses to evaluate the accuracy and scalability of the model.

VII. REFERENCES

- [1] López-Zambrano, J., L. J., & Romero, C. (2020). Towards Portability of Models for Predicting Students' Final Performance in University Courses Starting from Moodle Logs. *Applied Sciences*, pp.1-23.
- [2] Romero, C., & Ventura, S. (2020). Educational Data mining and Learning Analytics: An updated survey. *Data Mining and Knowledge Discovery*, pp.1-21.
- [3] Stracke, C. M. (2014). "Open learning for improving school education, lifelong learning and societal impact. The Fifth International Conference on e-Learning, (pp. 1-6). Belgrade, Serbia.
- [4] Shayan, P., & van Zaanen, M. (2019). Predicting Student Performance from Their Behavior in Learning Management Systems. *International Journal of Information and Education Technology*, pp.337-341.
- [5] Tempelaar, D. T., Rienties, B., & Giesbers, B. (2015). In search for the most informative data for feedback generation: Learning analytics in a data-rich context. *Computers in Human Behavior*, pp.157-167.
- [6] Gašević, D., Dawson, S., Rogers, T., & Gasevic, D. (2016). Learning analytics should not promote one size fits all: The effects of instructional conditions in predicting academic success. *The Internet and Higher Education*, pp.68-84.
- [7] Amrieh, E. A., Hamtini, T., & Aljarah, I. (2016). Mining Educational Data to Predict Student's academic Performance using Ensemble Methods. *International Journal of Database Theory and Application*, pp.119-136.
- [8] Kondo, N. O. (2017). Early Detection of At-Risk Students Using Machine Learning Based on LMS Log Data. 6th IIAI International Congress on Advanced Applied Informatics (pp. 198-201). Hamamatsu, Japan: IEEE.
- [9] Quinn, R. J., & Gray, G. (2020). Prediction of student academic performance using Moodle data from a Further Education setting. *Irish Journal of Technology Enhanced Learning*, pp.1-19.
- [10] Romero, C., Espejo, P. G., Zafra, A., Romero, J. R., & Ventura, S. (2013). Web usage mining for predicting final marks of students that use Moodle courses. *Computer Applications in Engineering Education*, pp.135-146.
- [11] Estacio, R. R., & Raga Jr, R. C. (2017). Analyzing students online learning behavior in blended courses using Moodle. *Asian Association of Open Universities Journal*, pp.52-68.
- [12] Wolff, A., Zdrahal, Z., Herrmannova, D., Kuzilek, J., & Hlosta, M. (2014). Developing predictive models for early detection of at-risk students on distance learning modules. In: *Machine Learning and Learning Analytics Workshop at the 4th International Conference on Learning Analytics and Knowledge*, (pp.24-28).
- [13] Pedregosa, F., Varoquaux, G., G. A., M. V., Thirion, B., Grisel, O. & Vanderplas, J. (2011). Scikit-learn: Machine learning in Python. *the Journal of machine Learning research*, 2825-2830.
- [14] Raga Jr, R. C., & Raga, J. D. (2017). Monitoring Class Activity and Predicting Student Performance Using Moodle Action Log Data. *International Journal of Computing Sciences Research*, 1-16.
- [15] Ramaswami, G., Susnjak, T., Mathrani, A., Lim, J., & Garcia, P. (2019). Using educational data mining techniques to increase the prediction accuracy of student academic performance. *Information and Learning Sciences*, 451-467.
- [16] Umer, R., Susnjak, T., Mathrani, A., & Suriadi, S. (2017). On predicting academic performance with process mining in learning analytics. *Journal of Research in Innovative Teaching & Learning*, 160-176.
- [17] Ferri, C., Hernández-Orallo, J., & Modroui. (2009). An experimental comparison of performance measures for classification. *Pattern Recognition Letters*, pp.27-38.
- [18] Dash, M., & Liu, H. (1997). Feature Selection for Classification. *Intelligent Data Analysis*, pp.131-156.
- [19] Uddin, M. F., Lee, J., Rizvi, S., & Hamada, S. (2018). Proposing enhanced feature engineering and a selection model for machine learning processes. *Applied Sciences*.
- [20] Ma, C., Yao, B, Ge, F., Pan, Y., & Guo, Y. (2017). Improving Prediction of Student Performance based on Multiple Feature Selection Approaches. In *Proceedings of the 2017 International Conference on E-Education, E-Business and E-Technology*, (pp. 36-41). Toronto.
- [21] Jović, A., Brkić, K., & Bogunović, N. (2015). A review of feature selection methods with applications. *International convention on information and communication technology, electronics and microelectronics* (pp. 1200 -1205). Opatija: IEEE.

Chapter 4

Perspectives on the challenges of generalizability, transparency and ethics in predictive learning analytics

The contents of this chapter are from the following article. In accordance with the Elsevier republishing policy, any full text that has been included, is the unmodified accepted article.

@2021 Elsevier. Reprinted, with permission, from: Mathrani, A., Susnjak, T., Ramaswami, G., & Barczak, A. (2021). Perspectives on the challenges of generalizability, transparency and ethics in predictive learning analytics. *Computers and Education Open*, 2, 100060.

<https://doi.org/10.1016/j.caeo.2021.100060>

The Elsevier open access policy permits the use, sharing, adaptation, distribution, and reproduction in any medium or format as long as appropriate credit goes to the authors.

Elsevier does not endorse any of Massey University's products or services. Internal or personal use of this material is permitted. If interested in reprinting/republishing Elsevier copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribution, please go to:


<https://www.elsevier.com/about/policies/copyright> to learn how to obtain a license from the correct link.

STATEMENT OF CONTRIBUTION DOCTORATE WITH PUBLICATIONS/MANUSCRIPTS

We, the candidate and the candidate's Primary Supervisor, certify that all co-authors have consented to their work being included in the thesis and they have accepted the candidate's contribution as indicated below in the *Statement of Originality*.

Name of candidate:	Gomathy Suganya Ramaswami
Name/title of Primary Supervisor:	Dr. Teo Susnjak
In which chapter is the manuscript /published work:	Chapter 4
Please select one of the following three options:	
<input checked="" type="radio"/> The manuscript/published work is published or in press <ul style="list-style-type: none"> • Please provide the full reference of the Research Output: Mathrani, A., Susnjak, T., Ramaswami, G., & Barczak, A. (2021). Perspectives on the challenges of generalizability, transparency and ethics in predictive learning analytics. <i>Computers and Education Open</i>, 2, 100060 	
<input type="radio"/> The manuscript is currently under review for publication – please indicate: <ul style="list-style-type: none"> • The name of the journal: • The percentage of the manuscript/published work that was contributed by the candidate: • Describe the contribution that the candidate has made to the manuscript/published work: 	
<input type="radio"/> It is intended that the manuscript will be published, but it has not yet been submitted to a journal	
Candidate's Signature:	Gomathy Ramaswami <small>Digitally signed by Gomathy Ramaswami Date: 2022.11.07 07:42:59 +13'00'</small>
Date:	07-Nov-2022
Primary Supervisor's Signature:	Teo Susnjak <small>Digitally signed by Teo Susnjak Date: 2022.11.07 15:30:10 +13'00'</small>
Date:	7-Nov-2022

This form should appear at the end of each thesis chapter/section/appendix submitted as a manuscript/publication or collected as an appendix at the end of the thesis.



Perspectives on the challenges of generalizability, transparency and ethics in predictive learning analytics

Anuradha Mathrani^{a,*}, Teo Susnjak^b, Gomathy Ramaswami^c, Andre Barczak^d

^a School of Mathematical and Computational Sciences, Massey University, Auckland, New Zealand

^b School of Mathematical and Computational Sciences, Massey University, Auckland, New Zealand

^c School of Mathematical and Computational Sciences, Massey University, Auckland, New Zealand

^d School of Mathematical and Computational Sciences, Massey University, Auckland, New Zealand

ARTICLE INFO

Keywords:

Learning analytics
Generalizability
Interpretability
Feature extraction
Transparency
Ethics protocol

ABSTRACT

Educational institutions need to formulate a well-established data-driven plan to get long-term value from their learning analytics (LA) strategy. By tracking learners' digital traces and measuring learners' performance, institutions can discern consequential learning trends via use of predictive models to enhance their instructional services. However, questions remain on how the proposed LA system is suitable, meaningful, and justifiable. In this concept paper, we examine generalizability and transparency of the internals of predictive models, alongside the ethical challenges in using learners' data for building predictive capabilities. Model generalizability or transferability is hindered by inadequate feature representation, small and imbalanced datasets, concept drift, and contextually un-related domains. Additional challenges relate to trustworthiness and social acceptance of these models since algorithmic-driven models are difficult to interpret by themselves. Further, ethical dilemmas are faced in engaging with learners' data while developing and deploying LA systems at an institutional level. We propose methodologies for apprehending these challenges by establishing efforts for managing transferability and transparency, and further assessing the ethical standing on justifiable use of the LA strategy. This study showcases underlying relationships that exist between constructs pertaining to learners' data and the predictive model. We suggest the use of appropriate evaluation techniques and setting up research ethics protocols, since without proper controls in place, the model outcome would not be portable, transferable, trustworthy, or admissible as a responsible outcome. This concept paper has theoretical and practical implications for future inquiry in the burgeoning field of learning analytics.

1. Introduction

The educational landscape has evolved with online learning management systems (LMSs) having facilitated any-time, any-place, and any-pace learning. Learners can interact with the different e-learning activities embedded in their institutional LMS; however, in doing so, they leave their digital traces or their digital footprints. Learner activities are captured via clickstream events associated with learners when they browse course content, navigate between course modules, download course material, participate over discussion forums, or upload/submit assignments (e.g., uploading a file for assessment or submitting an online quiz for marking). Learners' clickstream data annotated with the background LMS data are stored in log files [46, 50] that provide

digital footprint awareness to their educational institution. In other words, the digital footprint reveals experiences concerning "every article sound, image and information left, shared and clicked by the person [or learner] in the digital environment either consciously or unconsciously" ([49], p.50).

Learning analytics (LA) is concerned with sense-making of learners' digital footprints with the aim of understanding learner engagement patterns, such as, how learners traverse course structures and access course content at their own pace and time (e.g., How often do learners watch a video uploaded on the LMS by the teacher? Or, how often do they speed up, pause or rewind the video? [20]). By using such bottom-up approaches, LA can assist in creating new insights for further optimizing online learning experiences. LA employs educational data

* Corresponding author at: School of Mathematical and Computational Sciences, Massey University, East Precinct Gate 1, Albany Expressway (SH17), Albany 0632, Auckland, New Zealand.

E-mail address: a.s.mathrani@massey.ac.nz (A. Mathrani).

<https://doi.org/10.1016/j.caeo.2021.100060>

Received 15 February 2021; Received in revised form 11 November 2021; Accepted 14 November 2021

Available online 20 November 2021

2666-5573/© 2021 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

mining (EDM) methods (or reductionist techniques) for generating actionable insights which enable optimization of learning experiences. These methods consist of gathering clickstream data and combining them with other available learner data in order to generate meaningful features that describe learners' unique properties. Some of the commonly generated features include a number of learner logins into a LMS, their quiz completion times, grade averages and deviations from the cohort mean, number of course material downloads or the number of forum posts created/viewed amongst others. Once a rich set of features describing learners has been engineered on both current and past students, machine learning algorithms can then be used to generate predictive models. Subsequent steps frequently involve applying human judgement to the derived models so as to draw out insights which can facilitate enhancements to existing instructional services, and also inform institutions on future-oriented educational delivery strategies [47].

The field of LA is growing and there is ample evidence supporting its extensive applications in higher education [2,58]. However, many challenges still exist in conducting learning analytics and in achieving desired efficiencies. [50], p. 157] are of the view that "most institutions may not be ready to exploit the variety of available datasets for learning and teaching", since building a universal predictive model from the log files extracted from the LMS is not straightforward. The LMS data is scattered at different hierarchical levels that may be difficult to correlate [59]; moreover, the data have to be supplemented with additional data retrieved from other sources (e.g., student admission system, past study records). Once the combined data are pre-processed to a proper format, a predictive model is developed by inputting training datasets (or what is commonly referred to as *seen data*) to machine learning algorithms, whose outputs are then re-applied to new target data points (or to the *unseen data in respect to the algorithm*). Models which display high accuracy on unseen data are deemed to have generalized, or successfully transferred relevant knowledge from seen training datasets to the target data. In practice however, the training data and target data could differ in their composition and also the extent to which the target data represents the problem that is being posed which impacts the transferability or generalizability of the model. Further, ensuring interpretability of the model internals by providing explanations on the scope and rationale of the algorithmic functions used to generate them, is crucial for social acceptance of the models [52]. Watcher et al. suggest conveying human-understandable non-technical explanations to the intended learners on influences of specific features on the overall model predictions.

Alongside these issues also exist the ethical issues related to data privacy and data ownership. While institutions are privy to learners' course-related data, they also have access to learners' personal data (e.g., ethnicity, age, gender, prior study details, etc.), all of which must be used respectfully. The collection and usage of learner data by an institution has broader implications, such as, an increased power over the learner by their institution, or learners receiving little information on what data is being collected, or profiling learners based on race, socio-economic status, ethnicity or gender, all of which raise moral questions pertaining to intrusion on students' rights and privacy [27, 40]; hence, any data policy used for LA implementation must align with an institution's core principles [39] before it can be used for developing any form of institutional capability.

This section has briefly introduced some challenges commonly faced by a learning analytics enterprise. In the following section, we present four research questions that drive this study. Next we provide an overview of the state of science regarding current endeavors in establishing LA systems; since concept papers are about "what do we do, where have we come from, and what are the areas yet to be examined" rather than covering extensive literature reviews ([19], p. 128). In developing convincing arguments and providing theoretical explanations, concept papers assimilate and combine selected pieces of literary and empirical evidence to form a logical chain of argumentation [24]. This study

examines published literature that articulate implementation issues faced in LA, and in doing so, we direct the readers to key pieces of published literature that provide a deeper coverage of the major issues identified. Against this backdrop of previous studies and recent literature which explore general issues with the implementation of LA initiatives [31], we discuss operational challenges frequently encountered in building predictive models with educational datasets across multiple learner environments. Specifically, we discuss the generalizability, model transparency (which covers model interpretability and the explanation of predictions) as well as ethical concerns, and we suggest guiding frameworks to address them. Accordingly, a model generalizability framework is presented. The tensions faced in maintaining accuracy and effectiveness across low and high interpretability models are examined, and trade-offs around model transparency are identified. We further outline an institutional ethics protocol that can provide a regulatory structure for avoiding ongoing conflicts between having an algorithmic-driven strategy and maintaining learner privacy in a LA context. Finally, in the last section, we consolidate key points that emerged from our discussion to answer the research questions that can inform predictive model building activities for future analytics practice.

2. Research questions

While LA manifests as an innovative data-driven capability that can personalize learning based on individual learner needs, researchers need to evaluate the theoretical and methodological stance pertaining to the conduct of their analytics strategy. Researchers encounter non-trivial challenges at all stages of developing LA systems. These can be of a technical nature such as developing and selecting relevant features for predictive modeling, as well as making design choices about which type of predictive model to use given positives and negatives associated with different types. The difficulties can also be of a non-technical nature and concern how the predictive models are used and how their usage is communicated to the learners. This study therefore reflects on challenges associated with the development and deployment of LA systems to enable meaningful transformation of learners' data into relevant features that can lead to improved instructional services.

Following questions are posed.

- 1 What are the key challenges in effectively deploying LA systems?
- 2 What difficulties are still encountered in producing generalizable predictive models?
- 3 What are the next frontiers in being able to extract more value from predictive models, rather than just predictions?
- 4 Which ethical dilemmas still remain in the deployment and operationalization of LA systems?

3. Current studies on learning analytics

Learning analytics is envisioned by educational institutions as a powerful force that can lead to more personalized learner experiences. It is considered as a way to "track individual student engagement, attainment and progression in near-real time, flagging any potential issues to tutors or support staff" ([42], p. 6.). With the use of predictive models built from historical student datasets, many educational institutions have implemented strategies to boost student retention rates, maintain quality assurance practices, reveal key determinants for academic achievement, bring about self-regulated learning (by predicting individual learning needs) and enhance learners' experience [56]. But when a disconnect emerges between training datasets and target (live) data, the utility of the predictive models can be degraded [45]. Generalizability (or transferability) of the derived model [6] is constrained when the training and target data are extracted from different distributions that exhibit different learner scenarios. Researchers from educational fields, such as the EDM and LA community, are thus to some

degree restricted to the use of data belonging to similar courses when predicting students' performances.

In one study, the authors [9] built a universal predictive model from different MOOC (Massive Online Open Course) offerings. Data from three most recently finished MOOC course offerings, and also data from the initial weeks of an on-going course were used for building the predictive model. Using naïve algorithm and importance sampling approaches, they concluded that machine learning techniques should consider model performance on successive offerings of the same courses. The authors concluded that transferability can be improved when important sampling-based approach parameters are tuned by formulating a moving window size on longitudinal variables.

Successful transferability can take place in multiple ways, such as the reuse of some or all of the training data sets, or features extracted from those datasets. The transfer can also consist of reusing some model-specific settings extracted from a trained model to iteratively evaluate classifications in the target domain [23]. Hunt et al. put forward the transfer learning method for predicting students' graduation rates in undergraduate programmes. TrAdaBoost, an extended AdaBoost algorithm, was used to examine the effectiveness. That is instead of assuming all the training dataset (comprising a set of academic and demographic features of students belonging to different departments) came from the same distribution, the authors conducted two separate experiments each time using specific data for training. In the first experiment, the training set comprised all students apart from those studying engineering, while in the second one, the training set comprised all students that were suspended or on academic warnings. The experimental results showed that TrAdaBoost improved the accuracy of predictive models and recorded smallest error in both cases. Generally, TrAdaBoost helps improve the accuracy of predictive models by using the target set as a guide to select related data from the source set. However, when the target sample size is too small to be representative, TrAdaBoost does not improve performance because its selective process will be biased by the target samples and causes over-fitting to the target set. Moreover, when there is a variance in data distribution between training and target data, the predictive capability is compromised.

López-Zambrano, Lara, and Romero [32] proposed generic methods to check the feasibility of predictive models by grouping similar courses by degree or by similar level of usage of activities provided by LMS logs. Experimental results from a well-known classification algorithm (namely J48 from the Weka [55] software) showed that it is feasible to directly generate accurate models with an acceptable accuracy; however, the limitation is that the obtained models might result in low accuracy values with other courses that use different activities or actions compared to the course used for training. In such situations, the log files of the unseen data would show different events and then models efficacy would become compromised. [6]

Baesens, Ravi, Marsden, Vanthienen, and Zhao [5] add that deep analytic techniques (e.g., neural networks, support vector machines, ensemble methods) for building predictive capabilities rely on information (data) and trust, in which trust has not been given proper attention. Trustworthiness of any data-driven algorithmic decision relies on data quality and on fitness of data used that in turn inform specific learner features (e.g., number of forum posts or number of quiz attempts) on which the predictive capabilities are built. Researchers must have proper domain knowledge of the complex implementations of the underlying LMS and understanding of data characteristics in the digital footprint trail. This in turn will inform the internal design and improve predictability; and, position the LA enterprise in better providing human-understandable counterfactual explanations on the significance of any learner-related feature that has been considered by the model to impact the learners' performance [52]. With simple explanations, institutions can send the message across that they do not consider their learners as passive recipients. Explanations provide new grounds for conducting meaningful exchanges leading to ongoing interactions that further builds more trust in the operationalised predictive model.

"Building trust is essential to increase social acceptance of algorithmic decision-making" (p. 4); however, explaining the rationale and functionality of the algorithms that together computationally process the raw data with different rule-sets to provide predictor values is not an easy task. To appreciate the reductionist power of analytics and make sense of the predictor variables, learners must be able to comprehend how the predictions align with their digital footprint; therefore, interpretability of the models at a high level by its intended audience is crucial.

Rubel and Jones [40] raise further questions regarding the conflicting positions between student privacy and learning analytics. While they recognize the benefits of LA, they caution on usage of other forms of student personal data, such as the students' socioeconomic status, their demographic profile, academic history, or their financial aid package. Classifying data categories statistically based on socioeconomic status, race, or gender without proper thought could perpetuate "old prejudices" and "have a stifling effect on individuals and society" ([51]; p. 254). Proper controls that allow for differential access based on the merits of the purpose of data usage in learning analytics should be formed. That is, collecting learner information based upon their religious observances or politics amongst others would be impermissible under these controls; however, information based on learning patterns so as to nudge students for enhancing their learning outcomes would be endorsed. Tene and Polonetsky suggest disclosure statements be made by institutions on their usage of individual data that has been harvested in log files, but without disclosing the internal logic of their proprietary algorithms (which constitute their trade secrets). Further, some meaningful explanations should be offered on algorithmic interpretability so as to increase societal acceptance and build trust in the automated decisions from the predictive models [52].

This section has discussed some of the recent research works on LA and highlighted both the opportunities and limitations. In particular, we have emphasized on the generalizability, transparency of predictive models and ethical challenges faced as predictions are tailored across diverse course offerings and learner groups. Next, we propose some key points to address these issues.

4. Challenges in learning analytics

This section provides more perspectives to the challenges that have been identified in learning analytics literature. Concept papers, in particular, do not have data; rather their focus is on integration of domain concepts to offer propositions that can serve as a bridge between validation and usefulness [19, 53]. We highlight generalizability, model transparency and ethical domain challenges as some of the key areas from literature. Model generalizability refers to transfer learning issues with regard to the relevance of patterns extracted from the training datasets for use on new data. Educational institutions are accountable for ensuring that the model's predictive abilities are reliable, besides also holding a social responsibility of explaining the significance of model's predictions to their intended audience (or to their current learners), referred to as the model interpretability domain or explainable artificial intelligence. Moreover, the ethical domain related to the proper collection and usage of learner-related data is an important consideration for the LA enterprise. We discuss these challenges in more detail in the following three subsections.

4.1. Generalizability challenges

The intent of learning analytics is to uncover underlying relationship between predictors (e.g., assessment grades, participation via forum posts) and possible outcomes (e.g., final grades, course engagement level). A machine learning algorithm aids in computationally exploring data patterns in historic datasets and inferring rule-sets that can map predictor variables with outcomes being modelled. The result of this inductive process is a predictive model that can then be deployed to

make predictions on live data. However, achieving high predictive accuracies on real-world application data using these inductive methods is one of the key challenges. The failure of predictive models to generalize may happen for numerous reasons, and these will differ in respect to the unique challenges with which each application domain is associated. Machine learning challenges that are most relevant to the LA domain are “the curse of dimensionality”,¹ concept drift and class-imbalance ([41]. Moreover, the non-deterministic nature of the LA domain adds to the complexities around making accurate predictions about human behavior. While the dynamic nature of educational contexts can also compromise the predictive power induced from historic (training) data when applied to live (target) datasets. We observe this in dynamic environments when the training data used for deriving a predictive model ceases to correlate with the live data onto which the predictive model is being applied. In dynamic contexts such as LA, the underlying data may change frequently thus rendering the trained models using historic data, inaccurate. An example of this are models which have been developed for predicting learner outcomes for specific courses with strong dependencies on features representing different assessments. As the courses and the assessments evolve, or where the assessment syllabus and evaluation styles have changed, the historic training data risks losing relevance on the current live data. Such course-specific and highly tailored features are more powerful, but they have the potential to decrease the generalizability of the predictive models when their usage changes in subsequent deliveries [9, 16] .

Poor model generalization may also occur if the size of the training sample is not large enough for the machine learning algorithm to effectively create decision boundaries. In this instance high-dimensionality data are the culprit owing to the fact that it is tempting to exceed the number of features used in modeling in proportion to the size of the training datasets. Machine learning algorithms always uncover patterns. Many however are phantom patterns and do not correlate with reality. The more features an algorithm has access to, the higher the proclivity to discover meaningless patterns and overfit the model to irrelevant idiosyncrasies of the underlying dataset.

In the LA context, access to datasets used for machine learning may be limited to a few courses due to privacy and legal constraints in different jurisdictions, as well as to policies requiring opt-in consents from learners. Consequently, this may result in training datasets that are insufficient in size and therefore more prone to the negative effects arising from high-dimensionality data. More diverse and representative samples are critical for the field of LA research [16]. With small-sized datasets, the resulting models are even more likely to overfit by capturing residual noise rather than provide useful patterns. Alternatively, a model may underfit and thereby be unable to learn a complex decision boundary when the data volume is not rich enough to support this sufficiently. In either of the above scenarios, the accuracy obtained with the given training data may not match that of the models that have been deployed into a production environment, thereby limiting the usefulness of the derived models. Suggested ways to solve these issues are to make use of more general features, eliminating redundant or less-discriminatory features, incorporating more recent data points (while omitting some older data) and reducing the complexity of the models [11, 35]

It is well accepted within the machine learning community [12, 33] that the quality of features is more important than the choice of algorithms, or even the size of the training datasets. It has already been discussed that as courses evolve, often handcrafted features developed for an earlier course delivery may not express correlations with an ongoing course. Or they may not even exist in a subsequent course delivery.

¹ The curse of dimensionality is a phrase coined by Bellman [7], that refers to high-dimensional data, which in a LA context refers to learners' data that has a very large number of features or attributes describing each student. .

The problem of generalization is considerable. But this is further confounded by the phenomenon of *concept drift*. Concept drift refers to what happens to a predictive model over time as the training data and the current real-life data become disconnected. For example, prior to 2000s, virtual learning environments in LA were rare. These days, they are ubiquitous. This shift in technology also represents a gradual shift in the manner that students have come to learn and generate their digital footprints – and this shift continues today. The consequence is that using historic educational data that reaches too far back in time risks producing models which are not relevant for making decisions about current cohorts of students. However, using too little of the historic data for training also risks producing models that overfit. A difficult challenge arises in this domain where frequent concept drift needs to be accounted for and detected. How this can be accomplished, is still an active area of research [33].

In addition to the above, educational datasets tend to also be highly class-imbalanced. Class-imbalanced dataset domains possess an unequal representation in the number of samples for the different dependent variables. Classes with proportionally much fewer samples than the larger classes tend to experience a degradation in accuracy due to the fact that machine learning algorithms generally focus more on majority classes. As a result, predictive models may behave differently in terms of generalizability on majority and minority classes [44]. For instance, if the training dataset consists of very few students labelled as *at-risk* students (or students facing learning challenges), then there is an increased chance of misclassification in detecting these students on live data streams. The overall accuracy would likely be biased towards students not at-risk. The proposed solutions to handle class imbalances are pre-processing the dataset in order to construct a more balanced training dataset. This may be done by under sampling, over sampling or synthetic sampling methods [14]. Under sampling implies removing samples from the majority class (i.e., the not at-risk students) to balance with the minority class (i.e., the at-risk students); over sampling involves creating copies of the existing minority class samples (i.e., the at-risk students) to match the majority class (i.e., the not at-risk students); while synthetic sampling involves increasing the minority class with synthetic samples using feature space similarity. While strategies for mitigating class imbalances exist, it is not always clear which strategy should be used for a particular dataset and the overall challenge of machine learning on this type of data is also an active area of research with many open questions [26].

Another concern raised by [18] is that currently predictive models are very focused on transferring knowledge from the source domain (training dataset) to a target domain (live data) irrespective of whether these domains are related; thus potentially resulting in low generalizability. For example, consider an example of two sets of undergraduate students enrolled in a university. In this scenario, one set comprises final-year undergraduate students while the second set comprises first-year undergraduate students who have just entered tertiary study. These two sets represent different domains. The domains are obviously related (i.e., both belong to tertiary education), but their activity log files (extracted from the LMS) are likely to represent different learning patterns. The first-year undergraduates will interact with different activities within LMSs in a specific pattern due to having no previous experience with them. This will differ to the navigation patterns of the final year students who are more experienced. Moreover, the final year student cohort would have more background knowledge of their chosen area of study, hence their approach towards using online study resources would differ. These two domains are different and deriving a universal model between them could result in poor generalizability, or negative knowledge transfer. Fig. 1 below outlines these generalizability challenges as identified from literature.

The above challenges are highly relevant for the LA domain in respect to machine learning and generalizability. More broadly however, generalizability of machine learning models in the context of Big Data are also complicated by the presence of noisy data and the necessity

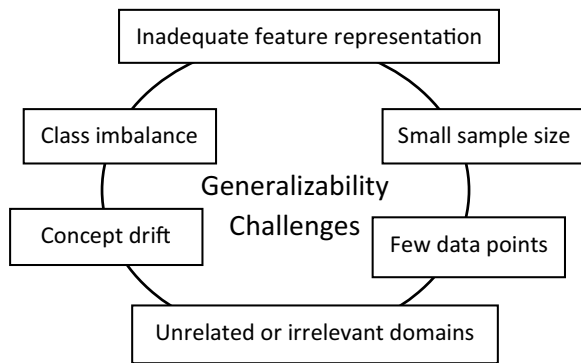


Fig. 1. Generalizability Challenges.

to learn with unreliable or contradicting data. This can also be compounded by the need to use sub-optimal algorithms due to the fact that the training data may be too large and cannot be processed and held entirely in a processing machine's working memory, which many algorithms require. These, as well as further challenges are reviewed in [28] and [57].

4.2. Model transparency (Interpretability and explainability) challenges

The rapid proliferation of predictive models into areas where previously human decision making was exclusive, has highlighted the need to be able to interrogate the mechanisms behind the models that drive their decisions. The goal is to ultimately generate glass-box models which provide transparency to the human-in-the-loop. There are a number of reasons why this is becoming important. Trust in black-box models is generally low, and trust in these systems can be forged through higher transparency. It is important to be able to verify the inner mechanics of the outputs of these algorithms in order to ensure that they are robust, reliable, and fair. Increasingly legal requirements are beginning to mandate that the predictive models account for their decisions and that the reasoning behind any automated decisions be clearly articulated to those affected by them. In addition, [52] point out the importance of those affected to have the ability to contest adverse decisions made by automated systems, and interestingly, to also have the ability to understand what would need to change in order to receive a desired result in the future, based on the current decision-making model.

The high interest in seeking new approaches to better understand the predictive modeling in real-life contexts such as education, has given rise to relatively new research fields such as Interpretable Machine Learning and Explainable Artificial Intelligence (XAI). The main goals of research in these spaces revolves around how *global model interpretability* and *model prediction explainability* can be achieved. Helpful literature surveys on these topics have emerged recently [1, 4, 13], together with some examples of some early work that is specific for the LA domain [3, 36]

Technically, *global model interpretability* deals with the challenge of making sense of the internals of a predictive model once a model has been trained by a machine learning algorithm. While *global model interpretability* highlights the behavior of the entire model at an abstract level, *model prediction explainability* on the other hand relates to the ability of a model to explain how it has arrived at a given prediction for a *specific* student. *Model interpretability* enables an institution to communicate to all students how a predictive model works using broad brushstrokes. *Model prediction explainability* enables an institution to respond to a specific student query about how and why they might have been identified as an at-risk student given this student's unique data.

Some algorithms produce models whose internals are in the form of decision trees or rule-sets which are highly interpretable at a global level. With these algorithms, it is easy to see the decision points and

threshold values for various features. However, higher accuracies are usually attained by algorithms that produce black-box models. Difficult trade-offs need to be made since some degree of accuracy or model interpretability will be sacrificed when choosing an algorithm. However, new suites of tools are emerging which are able to expose the internal logic of opaque models and induce them with adequate global interpretability, often through visualisations. Various approaches can be used such as generating proxy or surrogate models which approximate the underlying black-box model and generate interpretable models like decision trees (Trepan; [10]), rule-sets (BETA; [29]) or linear models (LIME; [37]). Apart from standard feature importance plots, more effective insights about the inner workings of models can be gleaned using tools that generate Partial Dependence Plots (PDP; [15]), which show how each feature affects the model's predictions across a range of values. While Individual Conditional Expectation (ICE; [21]) plots extend the PDPs with the ability to display the mean predicted outcomes for a range of values of a selected feature, meanwhile holding the values of other feature values constant. The challenge remains of matching the suitable tool for the particular educational dataset at-hand and performing extensive experimentation in order to identify the right tool.

In respect to explainability of predictions, global models with a high degree of interpretability can usually explain their individual decisions by highlighting the path through the decision tree that a single data point traverses, or in the case of rule-sets, listing the selected rules which were triggered by given predicates being met. In the instance of k -Nearest Neighbour models, k number of most similar students to the target student can be returned for inspection and comparison.

With opaque algorithms, once again additional tools are required in order to explain the model's reasoning. Recently, Shapley Additive Explanations (SHAP; [34]) have gained popularity in their effectiveness to visually explain the drivers of a model's decision-making process. Anchors [38] have also been recently developed as a tool that imparts a high degree of explainability. Anchors extend LIME by creating proxy models which are able to approximate non-linear functions and output a most succinct decision rule that "anchors" the prediction for a given data point for a given precision requirement. This means that rule anchors a prediction (the prediction will not change) with a given decision rule even if values change in other feature values, thus highlighting the key features for a given student. Using an opposite approach, counterfactuals [52] search out the smallest required change to a student's values which would result in a change of prediction. A counterfactual explanation in a case of a student, offers insights in terms of a minimum shift in key features that would need to take place in order to achieve a different outcome to what is currently predicted.

In summary, achieving full transparency and interpretability of operationalised predictive models in educational settings is challenging for a number of reasons already outlined, and presents delicate trade-offs that need to be made (refer Fig. 2). It can be tempting to use machine learning models that come with a high degree of intrinsic interpretability, but they produce less accurate models. The reverse is true with black-box algorithms. However, the trade-off is that additional tools need to be used in order to unpack both the internals of the models at an abstract level, as well as a suite of tools that provide explanations at an individual level of each student. The second scenario places additional burdens on educational providers to have larger and more skilled teams of data scientists who are able to work with a wide range of tools.

4.3. Ethical challenges

There is considerable evidence that confirms the value of learning analytics in the enhancement of institutional teaching environments. However, many institutions worldwide are still at early stages in their adoption of LA and in their practice of using data-informed approaches for improving instructional services and supporting learners [43], as they deal with associated ethical challenges. From an ethical standpoint, the field of learning analytics sits in contrast with other big data

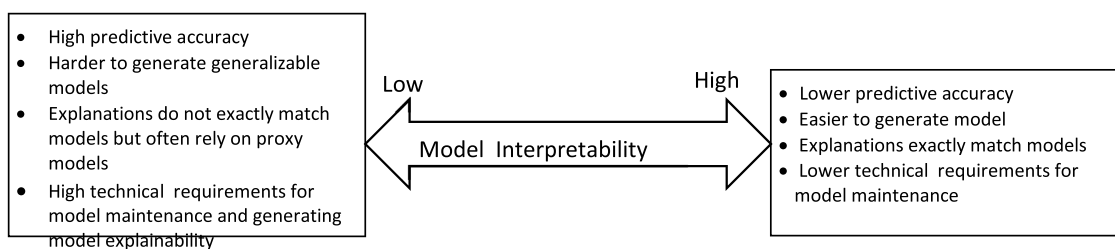


Fig. 2. Trade-offs between Low and High Model Transparency.

analytics (e.g., marketing analytics), since the digital footprints used are directly linked with individual students who can be identified via unique identifiers [40]. This raises questions about what constitutes acceptable or ethical analytics activities; that is, to what degree should the learners be informed about the details of how their data is used and whether explicit consent should be sought. Moreover, taking actions on automated predictions or recommendations from predictive models introduces levels of uncertainty, as future possibilities are conceived based upon their alignment with historic data. Institutions that use LA must confront lack of predictive certitude in deciding the effectiveness of predictive outcomes. Therefore, when flagging particular students, such as those who may be facing learning difficulties with intent to provide them with appropriate pedagogical interventions, they must provide learners with explanations on why such interventions are being actioned.

Interventions could comprise automated notifications that can nudge students by recommending relevant learning resources or by making some other provision in the form of personalized human assistance to help students overcome their learning difficulties [56]. Even though the intent of flagging students is to improve their learning environment, the fact remains that LA can also be considered to be a form of surveillance [25]. This tension between surveillance concerns and getting the true value out of LA has made it difficult to devise concrete ethical guidelines. Having said this, all scholarly research studies must follow high ethical standards. This involves conducting a proper ethical scrutiny by the concerned institution and by the analytics team to ensure *appropriate* protocols have been used in data collection and analysis. The word *appropriate* in the context of educational dataset for LA implies that institutional codes of conduct should cover elements of informed consent, privacy and de-identification, clearly state the scope and motive of learner data tracking, define the boundaries on data usage and have measures to prevent unauthorised access and disclosure of learner data [48]. However, getting ‘informed consent’ for participation would not be possible from past students (whose data has been used for training the model), or be feasible in covering large scale LA projects [27]; therefore, LA should be considered as “development or improvement of technological resources within an ethical framework” (p. 2862). Kitto and Knight further caution ethics committees, asking them to acknowledge the diversity across applied research disciplines from traditional education research, adding that “informed consent” may not always be possible within LA projects. Hence, we suggest definition of an institutional research ethics protocol that lays out detailed guidelines with respect to their technology deployment strategy that recognizes the purpose of the learner-generated data before leveraging any benefits from learning analytics.

Fortunately, the digital age has broadened everyone’s perspective on how digital footprints left on online public platforms can be leveraged by online agencies (e.g., advertising and marketing agencies). Online data traces can be linked to our persona such as to our social media profile, physical appearance, current location and to other personal interests, which can then be assessed by commercial agencies for their competitive advantage. Learners too are somewhat aware that they leave their digital footprints when interacting with the institutional LMS over the course of their study. However, if an institution intends to use

learners’ digital footprints for LA, they hold the responsibility of conveying their intention explicitly to the enrolled learners [48]. That is, they must reiterate to the learners that their online interactions are being recorded in the log files of a LMS; further, that the data from the user generated log files may be used by their institution for analytics. Therefore, as a first step, the research protocol should account for managing regulatory practices to ensure that learner privacy and confidentiality are not compromised when LA approaches are deployed for institutional advantage. In other words, institutions must acknowledge to all enrolled students that their digital footprints (captured via online interactions on LMSs) would serve as proxy data for analyzing their online behaviors. The proxy data would be mined and subsequently analyzed for gathering insights on learner behaviors that would in turn be used for improving overall instructional services. These services include creating models on user behavior, user experience, user profiles, trend analysis or for modeling various learner knowledge domains [8]. The benefits and limitations of these services must be explicitly stated in simple and non-technical language for ease in comprehension by the learners.

Moreover, in the case where historical learner datasets are to be used for developing instructional services, the provision of ‘informed consent’ from students no longer holds since these students are not currently enrolled at the institution. Therefore, institutions must ensure proper research protocols are followed to preserve the privacy and confidentiality of their past learner cohorts. First and foremost, instead of using actual unique student identifiers that can identify past students, the institution should follow a proper data management plan, such as to apply pseudonymization. Pseudonymization differs from anonymization where the “data subject is not or no longer identifiable”, since here the specific data subject cannot be identified without the use of “additional information” [17, 54]. The Article 4(5) adds that “such additional information is kept separately and is subject to technical and organizational measures to ensure that the personal data are not attributed to an identified or identifiable natural person”. In the context of LA implementation, pseudonymization techniques would imply replacing the real identifier of each student with another unique identifier that in no way can be connected to the actual student. Further, this procedure must be conducted in a fool-proof manner by institutionally approved data custodians to protect the re-identification of learners in compliance with Article 25. Hash algorithms (e.g., SHA-256) can be used to convert the unique student identifier into a fixed-length unique value that is used instead. The data stewards are therefore responsible that all personal identifiable information (e.g., name, address, and contact information) are removed and stored separately before using the pseudonymized data for model development. In this manner, the training data used for model development cannot be linked to any particular individual.

The key actionable output of LA systems are interventions. Since the aim is to essentially develop early warning systems that identify students who are at risk of underperformance or discontinuation, and to subsequently activate appropriate interventions in order to avoid these probable outcomes, clear understanding of what constitutes effective interventions must be known. However, the existing research into efficacy of various intervention types based on outputs of predictive models is unclear [30]. Further research needs to be conducted into effective

strategies for segmenting different learner types into groups which represent distinctive profiles, which ultimately have different needs and responsiveness to various types of intervention strategies. The instructional services henceforth produced from education data mining techniques will further inform on subsequent intervention strategies. Evidence-based strategies would relate to how current students who have been flagged as being at-risk or those facing learning difficulties are to be supported by tutoring staff in overcoming their learning challenges. However, there is danger of oversimplifying the intervention support strategy as an outcome of the model. We advise caution in simply setting up any intervention strategy and suggest that institutions consider multiple socio-pedagogical approaches for assisting students in overcoming their learning difficulties. Rubel and Jones [40] state that intervention strategies must stay clear of the student’s personal choices that are central to their conception of well-being or social tolerance (e.g., religion or politics). Instead, the strategies should be via tailoring of teaching practice or via interventions such as allocating relevant course resources conducive to improving student learning. Fig. 3 gives an overview of the ethical process that has just been described.

Meanwhile, more stringent legislation around data privacy like GDPR, require high levels of openness about operationalized analytics systems, and particularly the ability to explain to affected learners how certain automated decisions were formed, together with a list of all the contributing factors. While not all internals of a predictive model need to be explainable, there does however need to exist a mechanism that retraces prediction outputs for learners on-demand [22]. This is both a technical (Wachter et al., 2017) and a capability challenge which represents an ethical dilemma if predictive models are rushed into deployment without the ability to satisfy these requirements.

Finally, LA is a burgeoning field, and there is a dearth of educational datasets for the emergent researcher community to practice and hone their EDM skills. Another ethical concern faced by educational institutions is related to the sharing of student data with third parties in the current global environment [40]. Educational institutions are legally bound in ensuring privacy of student data thereby limiting

reproducibility and replication studies in learning sciences. The advent of MOOCs run by global providers (e.g., Coursera, edX, FutureLearn) offer another view of the emerging learning environments; although sharing of the learner data here too is restricted by strict privacy regulations [16]. Recognizing these restrictions, wherein full anonymity of each individual has to be maintained, many MOOC providers (e.g., HarvardX, Coursera) have released limited non-identifiable data via MORF (MOOC Replication Framework), a platform that allows researchers to deposit anonymized data (e.g., assessment details, grades, time stamps of student interactions, demographic information), that can be used by researchers in controlled environments while maintaining full privacy of student data. Further, researchers adhere to a global level ethics instrument that has been in place by the MOOC provider for responsible use of the anonymized data.

5. Conclusions, limitations and future directions

This paper has provided a much-needed perspective on the challenges encountered in deploying LA systems. Literature-based evidence in response to the first research question – What are the key challenges in effectively deploying LA systems? – has identified three challenges, namely, transferability or generalizability, model transparency, and ethical challenges. The LA movement espouses the premise that computational exploration of learners’ historical data could lead to relevant feature extraction and the development of predictive models that can profile currently enrolled students based on their learning needs. This can be further leveraged by the educational institution in facilitating intervention strategies for supporting learner communities in overcoming their learning difficulties. While it is tempting to have a general model solution that can be used across multiple courses and learner cohorts, in practice this is rather difficult to deliver. Moreover, model transparency concerns too need to be addressed for relevance, robustness, fairness, and social acceptance. Finally, the use of educational datasets has additional ethical concerns, such as responsible use of learners’ data so that individual learner’s privacy and confidentiality are

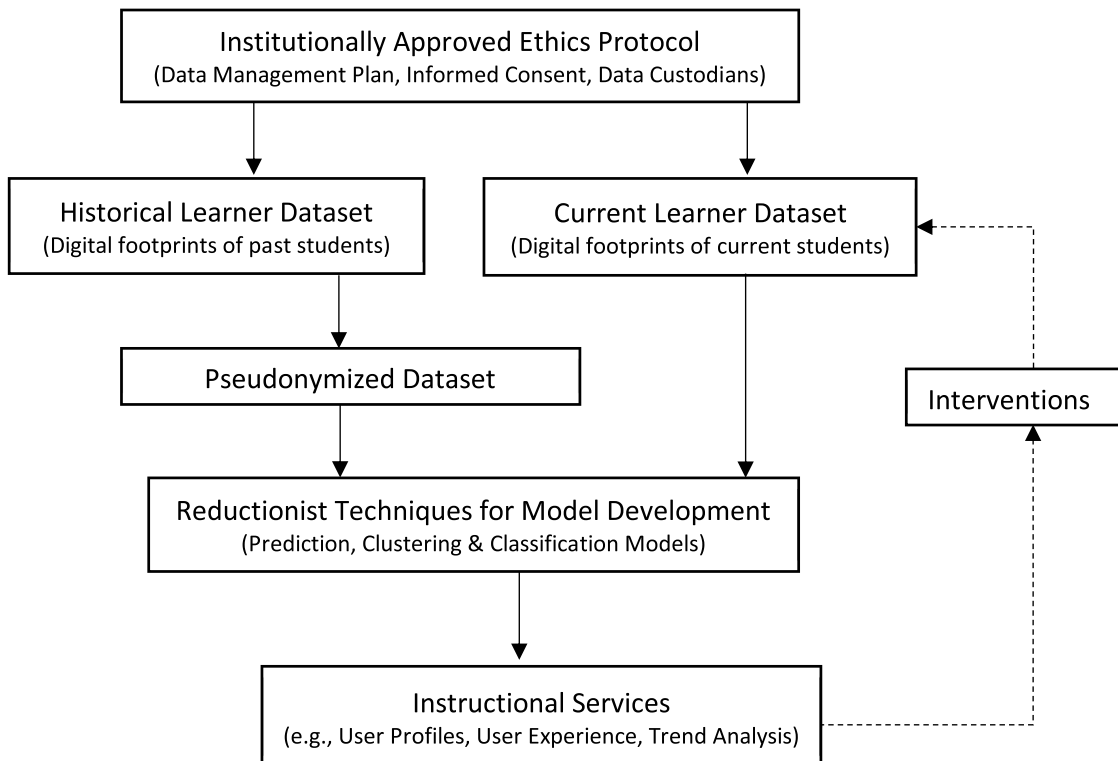


Fig. 3. Research Ethics Protocol.

not compromised.

Discussions pertaining to the second research question – What difficulties are still encountered in producing generalizable predictive models? – have revealed generalizability challenges associated with the dynamic nature of the domain, feature engineering and selection, small dataset sizes, unbalanced datasets, and concept drift amongst others (refer Section 4.1). For accurate transfer learning to take place by the prescribed models, educational institutions must first consider the constraints both in the use of learners' digital footprints and also the learning context. Exploratory analyses to acquire a basic understanding of the data and the learning context must precede any machine learning algorithmic analyses. Learning situations evolve with new course syllabus, changes to assessment structures, different learner cohorts and with diverse pedagogical approaches used by different tutoring staff. These lead to challenges in extracting relevant features that require much analysis and careful selection. However, a strong predictor in one learning scenario may become a weak predictor in another learning scenario. Therefore, we recommend re-evaluating the model design after regular intervals, such as after each successive course offering. Predictive models are driven by hindsight or historical data; it is crucial to ensure that the historical (or training) data aligns well with the target data to avoid generalization degradations brought on by concept drift. The analytics task is not a one-off task that concludes once a model is developed, rather it is an iterative empirical process of trial and error.

The third question – What are the next frontiers in being able to extract more value from predictive models, rather than just predictions? – has revealed gaps related to model transparency by the intended audience (refer Section 4.2). Algorithmic decisions as a consequence of predictive models are not easily interpretable. Hence, to achieve transparency, institutions must convey human-understandable explanations of the logic behind the internals of machine learning algorithms to their learners. In other words, simply informing a student about some algorithmic predictor value (e.g., AUC) as a measure of their learning behavior is by no means adequate; rather, both the high-level model behavior and the explanations of specific predictions must be made available in simple layman terms for non-technical people (i.e., to the currently enrolled students in this case). We have provided an overview of some strategies for explaining the model's reasoning by way of counterfactuals, proxy models and visuals (e.g., decision trees, feature importance plots, individual conditional expectation, etc.) as the next frontiers in extracting more value, rather than merely stating predictions.

The fourth question – Which ethical dilemmas still remain in the deployment and operationalisation of LA systems? – has further divulged ethical predicaments in the usage of learners' digital footprints being harvested within the institutional learning management platform. We acknowledge the crucial role learning analytics can play in transforming educational delivery with better flow of customized instructional services; however, we caution institutions on preserving the privacy and confidentiality of their learners. For any research to be recognized as a scholarly research outcome, all concerns related to its ethical conduct must be addressed first. However, we find that the ethical perspective in the deployment and operationalisation of LA systems is not explicitly stated in literature. We advise the use of an institutional research ethics protocol that clearly outlines the institutional strategy. Most importantly, disclosure statements on the use of learner data must be explicitly communicated to current learners, while historical data used for advancement/refinement of the model must be pseudonymized and all additional data that can lead to re-identification kept safe with authorized data custodians. The created model(s) is proprietary to the concerned institution; hence it is not required that institutions disclose their technical practices (e.g., ensemble of algorithms used).

This concept paper takes a problem-centered approach with the main purpose of “developing logical and complete arguments for associations rather than testing them empirically” ([19], p. 127); hence, no

experimental design or empirical data has been provided. Further, it does not cover analytic technicalities, such as choosing the right machine learning algorithm or tuning of the machine learning algorithms. Concept papers are meant to provide a tightly focused literature overview, since their objective is to put forth a bridge between validation and usefulness of constructs within some identified domain. Gilson and Goldberg advise the use of figures to clearly depict authors' views on how these constructs are related. Figs. 1, 2 and 3 showcase constructs for establishing and managing data-driven approaches related to the generalizability, model transparency, and ethical domains. While tracking and measuring learner performance can make education providers more aware of their instructional services, we encourage policy makers and institutional authorities to consider these constructs and question themselves on how their analytics approach is suitable, meaningful, and justifiable.

Declaration of Competing Interest

We have no affiliation with any organization with a direct or indirect financial interest in the subject matter discussed in the manuscript

References

- [1] Adadi A, Berrada M. Peeking Inside the Black-Box: a Survey on Explainable Artificial Intelligence (XAI). IEEE Access 2018;6:52138–60. <https://doi.org/10.1109/ACCESS.2018.2870052>.
- [2] Aldowah H, Al-Samarraie H, Fauzy WM. Educational data mining and learning analytics for 21st century higher education: a review and synthesis. Telematics Inf 2019;37:13–49. <https://doi.org/10.1016/j.tele.2019.01.007>.
- [3] Alonso JM, Casalino G. Explainable artificial intelligence for human-centric data analysis in virtual learning environments. 6-7 June 2019. Novedrate, Italy: Paper presented at the International Workshop on Higher Education Learning Methodologies and Technologies Online; 2019.
- [4] Angelov PP, Soares EA, Jiang R, Arnold NI, Atkinson PM. Explainable artificial intelligence: an analytical review. WIREs Data Mining Knowl Discov 2021;11(5): e1424. <https://doi.org/10.1002/widm.1424>.
- [5] Baesens B, Ravi B, Marsden JR, Vanhienen J, Zhao JL. Transformational issues of big data and analytics in networked business. MIS Q 2016;40(4):807–18. <https://doi.org/10.25300/MISQ/2016/40.4.03>.
- [6] Baker RS. Challenges for the Future of Educational Data Mining: the Baker learning analytics prizes. Journal of Educational Data Mining 2019;11(1b):1–17. <https://doi.org/10.5281/zenodo.3554745>.
- [7] Bellman R. Adaptive control processes: a guided tour. Princeton legacy library. Princeton, NJ: Princeton University Press; 1961.
- [8] Bienkowski MA, Feng M, Means B. Enhancing teaching and learning through educational data mining and learning analytics: an issue brief. Center for technology in learning. Washington, DC: U.S.: SRI International; 2012 (ED-04-CO-0040, <https://tech.ed.gov/wp-content/uploads/2014/03/edm-la-brief.pdf>, Task 0010).
- [9] Boyer, S., & Veeramachaneni, K. (2015, 2015/). Transfer learning for predictive models in massive open online courses. Paper presented at the Artificial Intelligence in Education, Cham.
- [10] Craven M, Shavlik J. Extracting Tree-Structured Representations of Trained Neural Networks. Adv Neural Inf Process Syst 1996;8:24030. <https://proceedings.neurips.cc/paper/1995/file/45f31d16b1058d586fc3be7207b58053-Paper.pdf>.
- [11] Ding, M., Wang, Y., Hemberg, E., & O'Reilly, U. (2019). Transfer Learning using Representation Learning in Massive Open Online Courses. Paper presented at the 9th International Conference on Learning Analytics & Knowledge, Madrid, Spain. <https://doi.org/10.1145/3303772.3303794>.
- [12] Domingos P. A few useful things to know about machine learning. Commun ACM 2012;55(10):78–87. <https://doi.org/10.1145/2347736.2347755>.
- [13] Doslilovic FK, Br j M, Hlupic N. Explainable artificial intelligence: a survey. In: The 41st International Convention on Information and Communication Technology, Electronics and Microelectronics. MIPRO; 2018. p. 0210–5.
- [14] Ebrahman SMA, Abraham A. A Review of Class Imbalance Problem. Journal of Network and Innovative Computing 2013;1:332–40. https://www.mirlabs.net/jnic/secured/Volume1-Issue1/Paper31/JNIC_Paper31.pdf.
- [15] Friedman JH. Greedy function approximation: a gradient boosting machine. Ann. Statist. 2001;29(5):1189–232. <https://doi.org/10.1214/aos/1013203451>.
- [16] Gardner J, Brooks C, Andres JM, Baker RS. MORF: a Framework for Predictive Modeling and Replication At Scale With Privacy-Restricted MOOC Data. Paper presented at. In: the 2018 IEEE International Conference on Big Data; 2018. <https://doi.org/10.1109/BigData.2018.8621874>.
- [17] GDPR. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (No. 02016R0679 - EN - 04.05.2016 - 000.002). Portal of the Publications Office of the EU; 2016. Retrieved from <http://data.europa.eu/eli/reg/2016/679/2016-05-04> Retrieved from <http://data.europa.eu/eli/reg/2016/679/2016-05-04>.

- [18] Ge L, Gao J, Ngo H, Li K, Zhang A. On handling negative transfer and imbalanced distributions in multiple source transfer learning. *Statistical Analysis and Data Mining: The ASA Data Science Journal* 2014;7(4):254–71. <https://doi.org/10.1002/sam.11217>.
- [19] Gilson LL, Goldberg CB. Editors' Comment: so, What Is a Conceptual Paper? *Group & Organization Management* 2015;40(2):127–30. <https://doi.org/10.1177/1059601115576425>.
- [20] Ginda M, Richey MC, Cousino M, Börner K. Visualizing learner engagement, performance, and trajectories to evaluate and optimize online course design. *PLoS ONE* 2019;14(5):e0215964. <https://doi.org/10.1371/journal.pone.0215964>.
- [21] Goldstein A, Kapelner A, Bleich J, Pitkin E. Peeking Inside the Black Box: visualizing Statistical Learning With Plots of Individual Conditional Expectation. *J Comput Graph Statist* 2015;24(1):44–65. <https://doi.org/10.1080/10618600.2014.907095>.
- [22] Holzinger A, Biemann C, Pattichis CS, Kell DB. What do we need to build explainable AI systems for the medical domain? *ArXiv preprint* 2017. <https://arxiv.org/abs/1712.09923>, (<https://arxiv.org/abs/1712.09923>).
- [23] Hunt XJ, Kabul IK, Silva J. Transfer learning for education data. 13–17 August 2017. Halifax, NS, Canada: KDD Workshop; 2017. Paper presented at the, <http://ml4ed.cc/attachments/HuntTransfer.pdf>.
- [24] Jaakkola E. Designing conceptual articles: four approaches. *AMS Review* 2020;10(1):18–26. <https://doi.org/10.1007/s13162-020-00161-0>.
- [25] Jones, K., & Salo, D. (2017). Learning analytics and the academic library: professional ethics commitments at a crossroads (No. id 2955779). Rochester, NY. Retrieved from <https://ssrn.com/abstract=2955779> Retrieved from <https://ssrn.com/abstract=2955779>.
- [26] Kaur H, Pannu HS, Malhi AK. A Systematic Review on Imbalanced Data Challenges in Machine Learning: applications and Solutions. *ACM Comput Surv* 2019;52(4): 79. <https://doi.org/10.1145/3343440>. Article.
- [27] Kitto K, Knight S. Practical ethics for building learning analytics. *British Journal of Educational Technology* 2019;50(6):2855–70. <https://doi.org/10.1111/bjjet.12868>.
- [28] L'Heureux A, Grolinger K, Elyamany HF, Capretz MAM. Machine Learning With Big Data: challenges and Approaches. *IEEE Access* 2017;5:7776–97. <https://doi.org/10.1109/ACCESS.2017.2696365>.
- [29] Lakkaraju H, Kamar E, Caruana R, Ileskovic J. Interpretable & Explorable Approximations of Black Box Models. *ArXiv preprint* 2017. <https://arxiv.org/abs/1707.01154>.
- [30] Larrabee Sønderlund A, Hughes E, Smith J. The efficacy of learning analytics interventions in higher education: a systematic review. *British Journal of Educational Technology* 2019;50(5):2594–618. <https://doi.org/10.1111/bjjet.12720>.
- [31] Leitner P, Ebner M, Ebner M. Learning Analytics Challenges to Overcome in Higher Education Institutions. In: Ifenthaler D, Mah D-K, Yau JY-K, editors. Utilizing learning analytics to support study success. Cham: Springer International Publishing; 2019. p. 91–104. https://doi.org/10.1007/978-3-319-64792-0_6. https://doi.org/10.1007/978-3-319-64792-0_6.
- [32] López-Zambrano J, Lara JA, Romero C. Towards Portability of Models for Predicting Students' Final Performance in University Courses Starting from Moodle Logs. *Appl Sci* 2020;10(1):354. <https://doi.org/10.3390/app10010354>.
- [33] Lu J, Liu A, Dong F, Gu F, Gama J, Zhang G. Learning under Concept Drift: a Review. *IEEE Trans Knowl Data Eng* 2018;31(12):2346–63. <https://doi.org/10.1109/TKDE.2018.2876857>.
- [34] Lundberg S, Lee SL. A unified approach to interpreting model predictions. *ArXiv preprint* 2017. <https://arxiv.org/pdf/1705.07874.pdf>, (<https://arxiv.org/pdf/1705.07874.pdf>).
- [35] Moubayed A, Injadat M, Nassif AB, Lutfiyya H, Shami A. E-Learning: challenges and Research Opportunities Using Machine Learning & Data Analytics. *IEEE Access* 2018;6:39117–38. <https://doi.org/10.1109/ACCESS.2018.2851790>.
- [36] Putnam V, Conati C. Exploring the need for explainable artificial intelligence (XAI) in intelligent tutoring systems (ITS). *IUI Workshops*; 2019. Paper presented at the.
- [37] Ribeiro, M.T., Singh, S., & Guestrin, C. (2016). "Why Should I Trust You?": explaining the Predictions of Any Classifier. Paper presented at the Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, California, USA. <https://doi.org/10.1145/2939672.2939778>.
- [38] Ribeiro, M.T., Singh, S., & Guestrin, C. (2018). Anchors: high-Precision Model-Agnostic Explanations. Proceedings of the AAAI Conference on Artificial Intelligence, 32(1) Retrieved from <https://ojs.aaai.org/index.php/AAAI/article/view/11491>.
- [39] Romero C, Ventura S. Educational data mining and learning analytics: an updated survey. *WIREs Data Mining and Knowledge Discovery* 2020;10(3):e1355. <https://doi.org/10.1002/widm.1355>.
- [40] Rubel A, Jones KML. Student privacy in learning analytics: an information ethics perspective. *Inf Soc* 2016;32(2):143–59. <https://doi.org/10.1080/01972243.2016.1130502>.
- [41] Sagi O, Rokach L. Ensemble learning: a survey. *WIREs Data Mining Knowl Discov* 2018;8(4):e1249. <https://doi.org/10.1002/widm.1249>.
- [42] Scater N, Mullan JS. Learning analytics and student success – assessing the evidence. *JISC* 2017. Retrieved 5/01/2021 Retrieved from <https://repository.jisc.ac.uk/id/eprint/6560> from <https://repository.jisc.ac.uk/id/eprint/6560>.
- [43] Scater N, Peasgood A, Mullan JS. Learning analytics in higher education: a review of UK and international practice. *VOCEDplus* 2016. Retrieved from www.jisc.ac.uk/reports/learning-analytics-in-higher-education, 39. www.jisc.ac.uk/reports/learning-analytics-in-higher-education.
- [44] Seah CW, Ong YS, Tsang IW. Combating negative transfer from predictive distribution differences. *IEEE Trans Cybern* 2013;43:1153–65. <https://doi.org/10.1109/TSMCB.2012.2225102>.
- [45] Shimodaira H. Improving predictive inference under covariate shift by weighting the log-likelihood function. *J Stat Plan Inference* 2000;227–44.
- [46] Siemens G. Learning Analytics: the Emergence of a Discipline. *Am Behav Sci* 2013; 57(10):1380–400. <https://doi.org/10.1177/0002764213498851>.
- [47] Siemens G, Baker RSJd. Learning analytics and educational data mining: towards communication and collaboration. Paper presented at the. In: Proceedings of the 2nd International Conference on Learning Analytics and Knowledge. British Columbia, Canada: Vancouver; 2012. <https://doi.org/10.1145/2330601.2330661>.
- [48] Slade S, Prinsloo P. Learning Analytics: ethical Issues and Dilemmas. *Am Behav Sci* 2013;57(10):1510–29. <https://doi.org/10.1177/0002764213479366>.
- [49] Surmelioglu Y, Seferoglu SS. An examination of digital footprint awareness and digital experiences of higher education students. *World J Educ Technol* 2019;11(1):48–64. <https://eric.ed.gov/?id=EJ1205431>.
- [50] Tempelaar DT, Rienties B, Giesbers B. In search for the most informative data for feedback generation: learning analytics in a data-rich context. *Comput Human Behav* 2015;47:157–67. <https://doi.org/10.1016/j.chb.2014.05.038>.
- [51] Tene O, Polonetsky J. Big data for all: privacy and user control in the age of analytics. *Northwestern J Technol Intellectual Property* 2013;11(5):239–73. <https://scholarlycommons.law.northwestern.edu/njtjip/vol11/iss5/1>.
- [52] Wachter S, Mittelstadt B, Russell C. Counterfactual explanations without opening the black box: automated decisions and the GDPR. *Harv J Law Technol* 2018;31(2): 841–87. <https://doi.org/10.2139/ssrn.3063289>. <https://ssrn.com/abstract=3063289> or <https://doi.org/>.
- [53] Weick KE. Theory Construction as Disciplined Imagination. *Acad Manag Rev* 1989; 14(4):516–31. <https://doi.org/10.2307/258556>.
- [54] Wes M. Looking to comply with GDPR? Here's a primer on anonymization and pseudonymization. *The Privacy Advisor* 2017. Retrieved from <https://iapp.org/news/a/looking-to-comply-with-gdpr-heres-a-primer-on-anonymization-and-pseudonymization/>.
- [55] Witten IH, Frank E, Trigg L, Hall M, Holmes G, Cunningham SJ. *Weka: practical machine learning tools and techniques with java implementations* (No. working paper 99/11). Hamilton, New Zealand: University of Waikato, Department of Computer Science. Retrieved from; 1999. <https://hdl.handle.net/10289/1040> Retrieved from <https://hdl.handle.net/10289/1040>.
- [56] Wong B, Li K, Choi SPM. Trends in learning analytics practices: a review of higher education institutions. *Interact. Technol. Smart Educ.* 2018;15:132–54. <https://doi.org/10.1108/ITSE-12-2017-0065>.
- [57] Zhou L, Pan S, Wang J, Vasilakos AV. Machine learning on big data: opportunities and challenges. *Neurocomputing* 2017;237:350–61. <https://doi.org/10.1016/j.neucom.2017.01.026>.
- [58] Ramaswami, G., Susnjak, T., Mathrani, A., Lim, J., & Garcia, P. (2019). Using educational data mining techniques to increase the prediction accuracy of student academic performance. *Information and Learning Sciences*, 120(7/8), 451–467. <https://doi.org/10.1108/ILS-03-2019-0017>.
- [59] Umer, R., Mathrani, A., Susnjak, T., & Lim, S. (2019, 30, March - 1 April 2019). *Mining Activity Log Data to Predict Student's Outcome in a Course*. Paper presented at the International Conference of Big Data and Education, London, UK. <https://doi.org/10.1145/3322134.3322140>.

Chapter 5

On Developing Generic Models for Predicting Student Outcomes in Educational Data Mining

The contents of this chapter are from the following article. In accordance with the Multidisciplinary Digital Publishing Institute (MDPI) open access policy, any full text that has been included, is the unmodified accepted article.

@2022 MDPI. Reprinted, with permission, from: Ramaswami, G., Susnjak, T., & Mathrani, A. (2022). On Developing Generic Models for Predicting Student Outcomes in Educational Data Mining. *Big Data and Cognitive Computing*, 6(1), 6. <https://doi.org/10.3390/bdcc6010006>.

The MDPI open access policy, permits the use, sharing, adaptation, distribution, and reproduction in any medium or format as long as appropriate credit goes to the authors.

MDPI does not endorse any of Massey University's products or services. Internal or personal use of this material is permitted. If interested in reprinting/republishing MDPI copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribution, please go to: <https://www.mdpi.com/authors/rights> copyright to learn how to obtain a license from the correct link.

STATEMENT OF CONTRIBUTION DOCTORATE WITH PUBLICATIONS/MANUSCRIPTS

We, the candidate and the candidate's Primary Supervisor, certify that all co-authors have consented to their work being included in the thesis and they have accepted the candidate's contribution as indicated below in the *Statement of Originality*.

Name of candidate:	Gomathy Suganya Ramaswami
Name/title of Primary Supervisor:	Dr. Teo Susnjak
In which chapter is the manuscript /published work:	Chapter 5
Please select one of the following three options:	
<input checked="" type="radio"/> The manuscript/published work is published or in press <ul style="list-style-type: none"> • Please provide the full reference of the Research Output: Ramaswami, G., Susnjak, T., & Mathrani, A. (2022). On Developing Generic Models for Predicting Student Outcomes in Educational Data Mining. <i>Big Data and Cognitive Computing</i>, 6(1), 6. https://doi.org/10.3390/bdcc6010006 	
<input type="radio"/> The manuscript is currently under review for publication – please indicate: <ul style="list-style-type: none"> • The name of the journal: • The percentage of the manuscript/published work that was contributed by the candidate: • Describe the contribution that the candidate has made to the manuscript/published work: 	
<input type="radio"/> It is intended that the manuscript will be published, but it has not yet been submitted to a journal	
Candidate's Signature:	Gomathy Ramaswami <small>Digitally signed by Gomathy Ramaswami Date: 2022.11.07 07:42:59 +13'00'</small>
Date:	07-Nov-2022
Primary Supervisor's Signature:	Teo Susnjak <small>Digitally signed by Teo Susnjak Date: 2022.11.07 15:30:10 +13'00'</small>
Date:	7-Nov-2022

This form should appear at the end of each thesis chapter/section/appendix submitted as a manuscript/publication or collected as an appendix at the end of the thesis.



Article

On Developing Generic Models for Predicting Student Outcomes in Educational Data Mining

Gomathy Ramaswami *, Teo Susnjak and Anuradha Mathrani

School of Mathematical and Computational Sciences, Massey University, Auckland 0632, New Zealand; T.Susnjak@massey.ac.nz (T.S.); A.S.Mathrani@massey.ac.nz (A.M.)

* Correspondence: g.ramaswami@massey.ac.nz

Abstract: Poor academic performance of students is a concern in the educational sector, especially if it leads to students being unable to meet minimum course requirements. However, with timely prediction of students' performance, educators can detect at-risk students, thereby enabling early interventions for supporting these students in overcoming their learning difficulties. However, the majority of studies have taken the approach of developing individual models that target a single course while developing prediction models. These models are tailored to specific attributes of each course amongst a very diverse set of possibilities. While this approach can yield accurate models in some instances, this strategy is associated with limitations. In many cases, overfitting can take place when course data is small or when new courses are devised. Additionally, maintaining a large suite of models per course is a significant overhead. This issue can be tackled by developing a generic and course-agnostic predictive model that captures more abstract patterns and is able to operate across all courses, irrespective of their differences. This study demonstrates how a generic predictive model can be developed that identifies at-risk students across a wide variety of courses. Experiments were conducted using a range of algorithms, with the generic model producing an effective accuracy. The findings showed that the CatBoost algorithm performed the best on our dataset across the F-measure, ROC (receiver operating characteristic) curve and AUC scores; therefore, it is an excellent candidate algorithm for providing solutions on this domain given its capabilities to seamlessly handle categorical and missing data, which is frequently a feature in educational datasets.

Keywords: machine learning; early prediction; CatBoost; at-risk students; educational data mining



Citation: Ramaswami, G.; Susnjak, T.; Mathrani, A. On Developing Generic Models for Predicting Student Outcomes in Educational Data Mining. *Big Data Cogn. Comput.* **2022**, *6*, 6. <https://doi.org/10.3390/bdcc6010006>

Academic Editor: Min Chen

Received: 23 November 2021

Accepted: 30 December 2021

Published: 7 January 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

In today's competitive world, higher education (HE) institutions need to deliver efficient and quality education to retain their students. The extensive blending of digital technologies into HE teaching and learning environments has resulted in large amounts of student data that can provide a longer-term picture of student learning behaviours. To understand and assist student learning, HE providers are implementing learning analytics (LA) systems; these systems provide data-driven insights about students to assist educators in determining their overall academic progression. This technology is particularly being leveraged to predict at-risk students and identify their learning problems at early stages, with the purpose of initiating timely interventions and tailoring education [1–3] to each student's level of need.

LA approaches typically rely on educational data collected from various learning activities provided via online learning platforms, such as learning management systems (LMSs) [4]. LMSs are web-based learning systems that offer a virtual platform for facilitating teaching and learning, such as providing students with online course content, tracking student interactions, enabling peer communication over online forums, delivering course assessments to students, or releasing assessment grades [5]. Various stakeholders have different objectives for using an LMS. For instance, Romero and Ventura [6] suggest that

students use an LMS to personalise their learning, such as review specific material or engage in relevant discussions as they prepare for their exams. Meanwhile, teachers rely on an LMS to deliver their course content and manage teaching resources in a relatively simple and uniform manner [7] without worrying about pace, place or space constraints. Irrespective of how an LMS is used, user interaction with the system generates significant and detailed digital footprints that can be mined using LA tools.

The most widely used open source LMS is Moodle (i.e., Modular Object-Oriented Dynamic Learning Environment). It facilitates instructors to create online lessons that can be used in any of the three delivery modes: face-to-face, entirely online, or blended [5]. Moodle is mostly used as a communication platform for educators to communicate with students, as well as to publish course materials and grade student assignments. Students can therefore benefit from Moodle by having more interactions with their instructor and their peers, as they engage with the course content. Every activity performed in Moodle is captured in a database or system log, which can then be analysed to examine underlying student learning behaviours via LA approaches. A deeper investigation may be conducted if any indicators pertaining to at-risk students are identified [7]. A modelling process translates these indicators (extracted from training data) into predictive insights, which can be used on new data (or test data) to gauge student online behaviours. Teachers can then support at-risk students in overcoming their learning difficulties [8].

Most of the research into predicting student performance and identifying at-risk students has focused on developing tailor-made models for different courses. There are multiple problems with this approach, such as the issue of scalability and overhead in developing, optimising and maintaining custom models for each course within the HE provider's vast array of course offerings. This approach brings to the fore both human resource and technical expenses. Even if these challenges can be overcome, course-specific models are likely to perform poorly across numerous courses due to data insufficiency, resulting in overfitting. This is a tendency for courses that may have small cohorts, or for courses that have been newly set up and thus do not have historic data from which patterns can be learned through machine learning.

An alternative approach is to develop generic predictive models that operate across all of the disparate courses, though only a few works have attempted this strategy to date [9]. This is a technically more challenging approach due to the fact that more experimental rigour is required in the feature engineering phase, whereby more generic and abstract features that describe students' learning patterns need to be devised in a way that they are course-agnostic. The present study focuses on building a generic (or portable) model using a machine learning algorithm to generate a classifier for predicting student final outcomes.

Therefore, the motivation of this study is to demonstrate how a generic, course-agnostic predictive model can be developed that has strong portability attributes and can thus be effective at predicting students' final outcomes across disparate courses. Furthermore, this research demonstrates how an effective model can be built using a variety of student attributes that range from demographic to those that capture students' academic performances. We perform numerous experiments by developing models at different time frames (two, four, six and eight weeks into a course) to examine how early in the course an effective prediction can be made. Our experiments involve multiple algorithms, and we note the CatBoost algorithm as being the most effective on our dataset.

Against this backdrop, the rest of the paper is organised as follows. The next section reviews existing studies on LA, with emphasis on the various machine learning techniques used to predict students' academic performance. Next, the datasets used in this study are described, followed by our proposed algorithmic approach. The results are then presented, findings are discussed, and conclusions are made thereafter.

2. Review of Related Studies

Prediction of student academic performance has drawn considerable attention in the educational field. For instance, predicting whether a student will pass or fail a course, and

then notifying the instructor about the at-risk student, can enable the instructor to intervene and provide the student with learning pathways to improve their performance [10]. Several studies have reported success in predicting student academic performance using various educational data mining techniques. This review covers related works on prediction using a generic (or course-agnostic) model and early prediction techniques within the scope of the study.

2.1. Related Study Exploring Prediction Using Generalised Model

Chen and Cui [11] applied a deep learning approach—long short-term memory (LSTM) networks—to analyse student online behaviours for early prediction of course performance using the students' LMS data. The prediction performance of the LSTM networks approach was compared against eight conventional machine learning classifiers and AUC was used as the evaluation metric. The model generalisability was evaluated using the data derived from semester 1 and semester 2. The results showed that test AUC scores for semester 1 were around 0.75, which was higher than those obtained for semester 2, as the training data was also from the course in semester 1.

The motivation of Zambrano et al. [9] was to study the portability of student performance where the knowledge extracted from a specific course can be applied directly to a different course from the same degree and with similar levels of LMS usage. Using J48 decision trees, the authors created a predictive model based on 24 courses to classify the students into a pass or fail category. The model achieved an AUC loss value of 0.09 and 0.28 when using the courses to the same degree, and an AUC loss range from 0.22 to 0.25 for courses with similar levels of Moodle usage.

LMS data and in-between assessment grades were used in another study [7] to predict student performance. Multiple linear regression was used to induce predictive models at the end of the course and evaluate the efficacy of the available features within LMS. To assess whether the data could offer an intervention during a course, linear and logistic regression was applied to the features at the end of each week of the course. The results showed the LMS usage data to be a weak predictor here unless other assessment data were also included, which ultimately improved predictions.

In another study by Nakayama et al. [12], the performance of students in a blended learning course was predicted based on their note-taking activities and their individual characteristics, which were measured via student surveys. The possibility of predicting performance in final exams was evaluated by using the features of the contents of notes taken by students throughout the course and overall participant characteristics. The results showed that features of note-taking activities play a major role in predicting the final exam scores.

Gasevic et al. [13] built different logistic regression models for nine undergraduate courses to predict student performance (pass or fail). They used LMS logs and student information from the institutional student information system to build one model that would cover all the courses, as well as one model per course. The authors computed the area under the ROC (receiving operating characteristic) curve (AUC) values. The generalised model for all the courses showed an acceptable accuracy ($0.5 \leq \text{AUC} < 0.7$). However, the models specifically built for a particular course achieved excellent ($0.8 \leq \text{AUC} < 0.9$) or outstanding ($\text{AUC} \geq 0.9$) performances.

2.2. Related Study Exploring Early Prediction of Student Performance

The main reason to predict student performance is to identify the at-risk students, intervening and customising learning strategies in time to support them in achieving better results; however, most earlier studies have focused on predicting students' final course results once all the student course data was gathered, which would leave no time for such interventions. Upon recognising this issue, more researchers have recently attempted to predict student outcomes earlier in their course of study. For instance, González et al. [14] tried to identify struggling learners at early stages of the course. They used LMS data to

predict student performance at set points along the way, when 10%, 25%, 33% and 50% of the course had been completed. Different classification algorithms, namely Decision Tree (DT), Naïve Bayes (NB), Logistic Regression (LR), Multilayer Perceptron (MLP), and Support Vector Machine (SVM), were used to evaluate the prediction accuracy. MLP achieved the best performance on this dataset, with 80% accuracy when 10% of the course had been delivered and 90% accuracy when half of the course had been completed.

A solution proposed by Queiroga et al. [14] used students' interaction with their virtual learning environment to identify at-risk students as early as possible. Classic DT, MLP, LR, random forest (RF), and the meta-algorithm AdaBoost (ADA) were used. The proposed approach utilised genetic algorithms (GA) to tune the hyperparameters of the classifiers, and the results were compared with the traditional method without hyperparameter optimisation. The prediction model was run every two weeks for a 50-week duration course. The results showed that the highest AUC score was achieved using the GA during the initial period. There was a considerable decrease in performance from week 30 due to the increase in the number of input attributes.

A study by Zhao et al. [15] used Moodle data to predict student performance in the first quarter of a semester. They used the fuzzy rule-based classification method (FRBCS) and a modified FRBCS to predict the learning outcomes of students. The results showed that the modified FRBCS method provided higher stability and better performance compared to the unmodified FRBCS method.

A study made by Ramaswami et al. [16] tried to estimate the earliest possible time within a course at which reliable identification of students at risk could be made. The input data were a combination of LMS, demographics and assignment grades. Four different classifier algorithms were tested: NB, RF, LR and k-Nearest Neighbours (kNN). Two experiments were conducted, one using all the features and the other using features selected for their high prediction accuracy. LR produced the best accuracy of 83% in week 11, and the authors noted that it is better to apply feature selection approaches rather than select all features for making predictions due to overfitting.

Howard et al. [17] attempted to determine the ideal time to apply an early warning system in a course. LMS data along with student grades and their demographic information were used as input data for the prediction model. After testing multiple predictive models, the Bayesian additive regressive trees (BART) model yielded the best results with a mean absolute error (MAE) of 6.5% as early as week 6, precisely midway through the course. This point in the course is sufficiently early so that remedial measures can be taken by the teacher, as required.

To sum up, various models have been proposed by researchers and a variety of different machine learning approaches have been used to mine educational data for student performance prediction. Moreover, some excellent results have been achieved using an assortment of different methods, with no one method outshining all other methods (refer to Table 1), unsurprisingly in line with the "No Free Lunch" theorem [18]. In general, a recent systematic literature review [10] into predicting student performance has found that reported accuracies range widely and are influenced by many factors, with the bulk of the studies appearing to achieve predictive accuracies between ~70% to ~90%.

Table 1. Machine learning algorithms applied in educational environments.

Authors	Prediction Goal	Evaluation Measures	Methods Compared	Best Performers
<i>Prediction Using Generalised Model</i>				
[11]	Binary classification	AUC	LSTM	LSTM
[9]	Binary classification	AUC, AUC loss	DT	Proposed method
[7]	Binary classification	Accuracy	Linear and logistic regression	Proposed method
[12]	Course grades	R-squares and prediction error	Support vector regression	Proposed method
[13]	Binary classification	AUC	LR	Proposed method
<i>Early Prediction of Students' Performance</i>				
[19]	Binary classification	AUC, F-measures	DT, NB, LR, MLP neural network, and SVM	MLP
[14]	Multiclass classification	AUC	DT, RF, MLP, LR, ADA, GA	GA
[15]	Binary classification	F-measures	FRBCS and modified FRBCS	modified FRBCS
[16]	Binary classification	F-measures, accuracy	kNN, RF, NB, and LR	LR
[17]	Final grades	MAE	RF, BART, PCR, KNN, NN, and SVM	BART

3. Datasets

The data used in this exploratory research was extracted from courses offered at an Australasian HE institution in a blended learning environment. Tayebinik and Puteh [20] have defined blended learning as a fusion of traditional face-to-face and online learning, where instructional delivery happens across both traditional and online courses, such that the online component becomes a natural extension of traditional learning. Data from various semesters were included in building the prediction model. The purpose of considering courses of different durations was to make sure our models could handle courses of different lengths and attributes. Single course-specific predictive models are commonly able to provide better performance when they have sufficient historic data, but they are not easily scalable to port to other courses. Data from the LMS (Moodle) action logs, the Student Management System (SMS) and the Enrolment Management System (EMS) were used for the study. These are described next.

3.1. Action Logs from Moodle

Moodle's built-in features track student activities in each course [21]. The courses comprise online modules pertaining to subject readings (provided via book resources, URL links or web pages), assessments (in the form of assignments or quizzes) and forum discussions. For the purposes of this study, the log data directly related to student activities were extracted from Moodle, while instructor data were excluded.

Each event record in the log signifies various actions (started, viewed, created, updated, etc.) performed by students on Moodle; data related to eight types of learning activities were collected. Table 2 shows the percentage of various activities logged by course. It should be noted that the usage of these activities varies across courses, and not all activities may be relevant for every course.

For instance, quiz activities do not form part of many courses, such as for Internet Programming or Application Software Development courses, amongst a few others; therefore, log data are not available for them (see Table 2). Numerical representations were calculated for each course module.

3.2. Enrolment Management System

The usage of the LMS and therefore the amount of logged data is scarcer in earlier parts of a course, which can be expected to lead to a lower predictive accuracy at initial stages of a course. Hence, demographic information and pre-academic data, such as age, gender, citizenship and entrance requirements, from the EMS were also utilised to augment the total dataset.

Table 2. Information about the courses and corresponding log attributes.

Course Name	Semester	Course Size		Number of Assessments	Grade Distribution		Logged Activities in %							
		Male	Female		High Risk	Low Risk	Forum	Quiz	Folder	Assign	Resource	Book	URL	Page
Introduction to finance	1	39	73	3	65	47	30	34.1	9.8	8.9	1.4	8.6	3.7	1.6
Introduction to finance	2	51	73	3	67	57	27	31.3	12.8	6.9	1.1	12.6	4.6	2.5
Computer Applications and the Information Age	1	67	48	4	94	21	42.1	-	1.1	24	21.1	7	0.9	2.9
Computer Applications and the Information Age	3	21	12	4	20	13	43.5	-	0.8	23.9	22.1	5	0.9	2.9
Fundamentals of Information Technology	1	78	15	3	68	25	34.2	-	-	23.1	17.4	4.9	16.6	3.1
Fundamentals of Information Technology	2	75	42	6	71	46	35.8	-	-	22.2	16.6	5.8	17.2	1.5
Application Software Development	1	87	10	6	58	39	23.4	-	-	34.7	34.3	-	7.5	-
Internet Programming	2	61	5	4	37	29	33.7	-	0.3	17	37.1	-	7.3	4.2
System Analysis and Modelling	2	68	25	4	65	28	33.1	7	4.1	8.4	21.2	16.5	6.6	0.09

3.3. Student Management System

The SMS provided assessment data consisting of in-between assignment grades, quizzes, final exams and total course marks. Some of the assessments were online and logged in the Moodle LMS, while other assessments were offline and handed in on paper or via other systems.

4. Method

Before the students' data could be accessed for conducting any form of analysis, ethics approval was required. An ethics application was submitted to the Human Research Ethics Committee at the host HE institution, and subsequently, approval to proceed was attained.

Once the data were obtained, a design decision was made to use a derivative of the final course mark for each student as the target (or dependent) variable for prediction. The final course mark is based on weighted averages of the marks that students received from the online assignments and the final examination.

Student support services at the given institution require predictive outputs that indicate only two possible outcomes for each student, namely if they are at risk or otherwise. Therefore, in order to conduct machine learning, each student in the historic dataset record needed to be labelled as being either at risk or not. A subsequent design decision needed to be made on how to define each of these two categories so that models could be generated. An early assumption was made that students with a mean final course mark of 50% or less were likely to be at-risk students due to the fact that these students have not successfully completed numerous courses. However, initial experiments using this threshold yielded highly imbalanced datasets and consequently poor models.

Subsequently, the threshold for defining at-risk students was adjusted and fixed at 60% or less for a mean final course mark, with the remainder being considered not at-risk. This threshold decision was supported by the underlying data, which revealed that the mean course mark for students who eventually abandoned their qualification studies was 61%, and only 15% of students who eventually completed their qualification studies achieved a mean course mark of 60% or less. The adjusted threshold addressed the class imbalance problem to a sufficient degree since most machine learning algorithms can handle some imbalance, and standard performance evaluation measures can still be effective unless large amounts of imbalance exist [22].

The raw data from the LMS log files were processed in order to engineer features used for building the predictive models. This data processing step converted raw data into variables that captured normalised and relative attributes of each student in respect to their cohort. In doing so, features were generated that were course-agnostic and thus generic. This entailed calculating the rolling mean of the actions performed by students over multiple weeks rather than displaying the counts of the actions performed by students each week. The rolling mean was calculated from the averages of actions from every week and is represented as a single column, hence reducing the number of feature columns. In addition to the rolling mean, the Z-score (or standard score) was also calculated, which relativised a student's score on a given feature with that of their peer's based on the degree of deviation from the mean. When in-between assessment grades were available, they were added to the input data and the same procedure was followed, which meant that the assessment features were not tightly coupled with the specific courses.

$$Z = \frac{X - \mu}{\sigma} \quad (1)$$

In Equation (1), X denotes the value of the independent variable, μ is the cohort mean score for the independent variable and σ is the standard deviation of the independent variable.

Students' prior course grades had an impact on the students' performance, as the prior course grade was linearly related to the final score (Figure 1). Hence, the prior grades of the students along with the count of pass/fail for previous courses were measured. Table 3 represents the various numerical and categorical features that were used for this study.




Figure 1. Comparison of prior score with final score of students.

Table 3. Feature description.

Feature Name	Description	Type
Average score of prior courses	The mean score achieved by a student from across all previous course scores	Numerical
Maximum score achieved in prior course	The maximum score achieved by a student from their previous courses	Numerical
Prior course deviation score	The Z-score of a student in respect to the deviation of the cohort mean	Numerical
Assignment score	The assignment scores received by a student	Numerical
Assignment deviation score	The Z-score of the student's mean assignment score as a deviation from the cohort mean	Numerical
Prior role description	Student's previous year's primary activity	Numerical
LMS deviation score	The engagement score expressed as a Z-score of a student as a deviation from the cohort mean.	Numerical
LMS engagement score	The count of all activities performed by a student on the Moodle platform.	Numerical
Citizenship	The nationality of the student	Categorical
Age	Age of a person	Categorical
Highest school qualification	Highest school qualification at admission	Categorical
Study mode	Study by distance/online or on-campus	Categorical
Gender	Gender of the student	Categorical
English proficiency test	English proficiency	Categorical

4.1. Predictive Modelling

The prediction was performed using classification algorithms that classify a given instance into a set of discrete categories; in our case, the pre-defined categories are at risk or not at risk. There is a wide range of algorithms supporting classification used throughout the literature. Choosing the optimal algorithm is difficult since they differ in numerous aspects, such as learning rate, robustness and the amount of data required for training, as well as their biases and behaviours on different datasets.

This study used the recently developed CatBoost [23] algorithm for prediction tasks. CatBoost was chosen since it is flexible in its ability to work with both categorical and numeric features and is able to seamlessly function in the presence of missing values. A categorical feature is a feature that has a discrete set of values that are not essentially comparable with each other. In practice, categorical features are usually converted to numerical values before training, which can affect the efficiency and performance of the algorithm. However, CatBoost has been designed with the specific aim of handling categorical data. The ability to handle both categorical and missing data represents a considerable technical advantage over other algorithms, and a recent interdisciplinary review [24] has found it to be competitive with other state-of-the-art algorithms.

CatBoost is a category boosting ensemble machine learning algorithm that uses the gradient boosting technique by combining a number of weak learners to form a strong learner. It does not use binary substitution of categorical values; instead, it performs a random permutation of the dataset and calculates the average label value for every object.

The combinations in CatBoost are created by combining all categorical features already used for previous splits in the current tree with all categorical features in the dataset. CatBoost thereby reduces overfitting, which leads to more generalised models [23]. CatBoost uses ordered target statistics (Ordered TS) to tackle categorical features for a given value of the categorical feature, which means that the categorical feature is ranked before the sample is changed with the expectation of the original feature value. In addition, the priority and its weights are included. In this way, the categorical features are changed into numerical features, which effectively decreases the noise of low-frequency categorical features and improves the robustness of the algorithm [25]. The performance of CatBoost was benchmarked against four other standard machine learning algorithms: Random Forest [26], Naïve Bayes [27], Logistic Regression [28] and kNearest Neighbours [29].

4.2. Machine Learning Procedure

The machine learning procedure used the hold-out method to perform the experiments where the order of occurrences of the training-test split was maintained. The experiments were implemented using Python [30], which has built-in functions suitable for estimating and refining the results of predictions.

4.3. Model Evaluation

After designing the classification model, the next step is to evaluate the effectiveness of the model. This study used total accuracy, recall, precision and the F-measures derived from a confusion matrix [31]. The total accuracy alone can be a misleading evaluation measure if the dataset is imbalanced because a model often favours learning and the prediction of a value of the most frequent class. This can give a misleading impression that the classifier has generalised better than it really has. In such circumstances, it is preferable to use the F-measure, which considers both precision and recall; an F-measure is the harmonic mean of precision (or positive predictive value) and recall (sensitivity) shown in Equation (2).

$$F - measures = 2 * (Precision * Recall) / (Precision + Recall) \quad (2)$$

Precision is the measure of classifier's exactness

$$Precision = \frac{TP}{TP + FP} \quad (3)$$

Recall is the measure of the classifier's correctness.

$$Recal \text{ or True Positive Rate} = \frac{TP}{TP + FN} \quad (4)$$

where *TP* denotes TruePositive, *FP* signifies FalsePositive, *FN* represents FalseNegative, and *TN* denotes TrueNegative.

Further, the AUC (Area Under the Curve)—ROC (Receiver Operating Characteristics) curve is a performance measurement for a classification problem at several threshold settings [32]. ROC is a probability curve, and AUC signifies a measure of separability and is used for distinguishing between the class labels. AUC rates range from 0 to 1. The acceptable AUC range for a predictive model depends on the context. Typically, in most research areas, an AUC rate above 0.7 is preferred [33]. In general, higher AUC scores indicate models of higher quality. The following equations are used for calculating AUC-ROC. In our study, we used the accuracy, F-measure and AUC as evaluation metrics. Our decision to use accuracy, F-measure and AUC as evaluation metrics is supported by a recent systematic literature review [10] into the usage of predictive models in LA contexts.

4.4. Experimental Design

Our experimental design was devised in a manner that robustly evaluated the ability of our generic models to port across different courses and different deliveries of the same courses. We used a modified k-fold cross validation approach to evaluate our generic models. Given that our dataset is made up of seven different courses, and a total of 10 separate deliveries of those courses, we decided to train a model using nine course deliveries and test against the remaining hold-out course offering. We repeated this process 10 times with a different combination of training and hold-out courses in order to arrive at our final, aggregated evaluation scores for our models.

5. Results

The results from our evaluations in regard to portability of the generic (or course-agnostic) prediction model and feature selection approaches are described here.

5.1. Overall Performance Comparison of Classifiers

The generalisability of all the generic classifier models is shown in Table 4, which records the aggregated model evaluation across all the hold-out datasets, together with the standard deviation over the 10 separate train/test executions. The best performing scores are in bold.

Table 4. Performance scores of various classifiers.

Classifiers	F-Measure	Accuracy	AUC
CatBoost	0.77 ± 0.024	75 ± 2.1	0.87 ± 0.023
Random Forest	0.67 ± 0.025	67 ± 2.4	0.74 ± 0.015
Naïve Bayes	0.67 ± 0.023	68 ± 2.3	0.71 ± 0.034
Logistic Regression	0.68 ± 0.031	67 ± 3	0.73 ± 0.025
K-Nearest Neighbors	0.71 ± 0.02	71 ± 2.4	0.72 ± 0.022

We can observe from the results that the F-measure scores range from 0.67 to 0.77, with CatBoost achieving the highest score. The variability across the different algorithms on all 10 datasets is similar, indicating that all algorithms generated generic models with a stable behaviour on all the courses.

The overall accuracy of all the algorithms has similarly ranged between 65% and 75%, with the result from CatBoost again clearly outperforming the other algorithms on these datasets. The accuracies of the bottom three algorithms, namely Naïve Bayes, Random Forest and Logistic Regression, did not exhibit significantly divergent accuracy results.

The AUC is a particularly effective measure of the effectiveness of the decision boundary generated by the classifiers to separate out the two class labels representing at-risk and not at-risk students. These scores range from 0.71 to a very high score of 0.87, with CatBoost once more producing the best result and demonstrating that it has generated a classifier with the most effective decision boundary.

5.2. Classifier Performance Snapshots

The experimental results indicate that CatBoost is the most effective algorithm on these datasets and feature set, across the various standard algorithms considered in the study. Our attention now turns to examining the behaviour of the best performing algorithm on one of the hold-out datasets (Computer Applications and the Information Age—Semester 3). Our motivation was to analyse the ability of the algorithm to produce generic models that can identify at-risk students at early stages of a course so that in a practical setting, timely interventions can be initiated. To that end, we partitioned the dataset into 2, 4, 6 and 8 week intervals. We trained a CatBoost classifier on each snapshot/partition of the dataset in order to simulate a real-world scenario where only partial information is available at different stages and points in time of a course.

Figure 2 depicts the F-measure and AUC score for the CatBoost classifier on the hold-out dataset, depicting the metric scores at each of the 2, 4, 6 and 8 week snapshots in time. As expected, we observe that there is generally an improvement in the generalisability of the classifier as a course progresses and more data and a richer digital footprint are acquired for each student. However, the predictive accuracy as expressed by both AUC and the F-measure is still very high at an early 2-week mark, meaning that an accurate identification of at-risk students can already be made within two weeks of a course's commencement, and that necessary interventions can be conducted as required in a timely fashion. Similar patterns were observed on the remaining nine hold-out datasets at the same snapshots in time.




Figure 2. F-measure and AUC plot of CatBoost on the Computer Applications and the Information Age—Semester 03 hold-out dataset.

5.3. Feature Importance

Feature importance analysis is an important aspect of machine learning, as it enables practitioners to understand what the key factors are that drive the predictions. In this analysis, we are able to quantify how much each feature of the data contributes to the model's final prediction, thus introducing some measure of interpretability.

In this study, we used the Shapely Additive Explanations (SHAP) method [34] for estimating feature importance and model behaviour. The SHAP method has the additional ability to depict how changing values of any given feature affect the final prediction. The SHAP method constructs an additive interpretation model based on the Shapley value. The Shapley value measures the marginal contribution of each feature to the entire cooperation. When a new feature is added to the model, the marginal contribution of the feature can be calculated with different feature permutations through SHAP. The feature importance ranking plot from the training data is shown in Figure 3, representing the generic model from Section 5.1.

In the SHAP feature importance graph seen in Figure 3, each row signifies a feature. The features are arranged from the most important at the top to the least consequential at the bottom. The abscissa corresponds to the SHAP value (Figure 4), which influences the final prediction. Every point in the plot denotes a sample, where red represents a higher feature value and blue represents a lower feature value. The vertical line on the plot, centred at 0 represents a neutral contribution towards a final prediction. As the points on the graph move further to the right on the x -axis from this vertical line, the higher the positive contribution becomes towards the prediction of a student succeeding. The inverse is true for the strength of a contribution towards a prediction for an at-risk student as the values extend into the negatives on the x -axis.

In our model, the assignment grades (both current and prior grades) played a vital part in prediction, which is unsurprising. We can observe from the graph that as the maximum grade for a student increased, the stronger this became as a predictor for positive outcomes. The reverse relationship also held; however, it was more acute, indicating that lower assignment grades had a stronger negative contribution than higher assignment grades for positive outcomes.




Figure 3. Feature importance ranking plot with CatBoost SHAP.




Figure 4. Shapley force plot of a sample.

The above figure depicts the behaviours of the model at a global level. However, the SHAP method can also explain a given prediction for an individual student using a SHAP force plot (refer to Figure 4). The force plot depicts the extent to which each of the most significant features is pushing a final prediction towards a negative or positive prediction. The ‘base value’ on the graph represents the mean prediction value, which can be regarded as a 50/50 point. Values to the right of the base value represent a final positive prediction and the inverse for the final values on the left. In the example below, we see that the assignment deviation score, the rolling assignment average score and the part-time study mode of a student most strongly influence the final prediction towards a positive prognosis for this specific student. Meanwhile, the age of the student has some influence towards a negative outcome prediction.

While feature importance shows what variables affect predictions the most, and force plots indicate the explainability of each model’s predictions for an individual student, it is also insightful to explore interaction effects between different features on the predicted outcome of a model. A dependence scatter plot seen in Figure 5 demonstrates this. The *x*-axis denotes the value of a target feature, and the *y*-axis is the SHAP value for that feature, which relates directly to the effect on the final prediction. Figure 5a denotes that there is generally a positive correlation between LMS engagement scores and the assignment

scores. Two noteworthy patterns emerge from Figure 5a. First, we can see that students who score highly on the assignment scores but exhibit poor engagement levels with the LMS receive strongly negative outcome predictions by the model. We also see there is a threshold of 0.4 for the LMS engagement score, and those that score above this threshold are positively correlated with the model's predictions for successful outcomes, which is amplified further for those with higher assignment scores.




Figure 5. Two SHAP scatter dependence plots: (a) left graph depicting the interaction between LMS engagement scores and assignment scores, (b) right graph depicting the interaction between learner age and LMS engagement scores.

Figure 5b considers the feature interaction of learner age and LMS engagement scores. Age of a student plays a significant role in the generated models. As the age of a student increases, the effect on the model predictions for positive outcomes becomes stronger. From approximately age 26 onwards, increases in the student age carry a stronger positive effect on model predictions for positive outcomes until age 40, from which point there do not appear to be any further increasing positive effects. Higher LMS engagement scores generally interact positively with age for successful outcomes. It is striking that those who are most at-risk are those in their early twenties with LMS engagement scores having no clear positive effect on prediction outcomes for this student demographic.

6. Discussion

The results have clearly indicated that generating viable generic models that are not tightly coupled to specific attributes of different courses is possible. The generic model has shown an effective average accuracy of around 75% across quite diverse courses. The diversity was seen in the fact that courses had varying numbers of assessments, differing distributions in their usage of the messaging Fora. Some courses used quizzes, books and various online learning resources, while others did not, or they differed greatly in their emphasis of their usage. The generic models were able to handle widely disparate courses, and despite them, produce useful predictive models even at earliest stages of a course's delivery in order to permit timely interventions for at-risk students if necessary.

The 75% accuracy was stable across all the diverse courses, and it is comparable to predictive accuracies attained in published research. It could be argued that 75% accuracy may not be high enough to instil sufficient confidence in the models. However, one needs to keep in mind that the defined categories of 'at-risk' and 'not at-risk' are not black-and-white, clear-cut categories. These categories are moving targets and students are likely to fall into either category at different times of their study, or perhaps during different courses that they might be undertaking. Therefore, these categories embody many grey areas, and

due to the lack of hard boundaries, the predictive accuracy can always be expected to be somewhat compromised when converting this complex problem describing a shifting continuum into a well-defined binary problem. One strategy to overcome this limitation is to instead focus on success prediction probabilities, which produce continuous valued outputs between 0 probability of success to 1, denoting complete confidence in successful outcomes. In addition to relying more on these outputs rather than on hard-thresholded binary categories, one can produce weekly probability predictions and monitor the change in the delta from one week to the next in order to generate deeper insights into the risk profile of a given student.

As for the superior performance of CatBoost over the other algorithms on these datasets, some of this can be attributed to the higher complexity of CatBoost over the other algorithms used in this study. CatBoost is a gradient boosting algorithm that makes it considerably more sophisticated than our benchmarking algorithms, capable of inducing more complex decision boundaries with some safeguards against overfitting. These internal mechanics undoubtedly contributed, but also the fact that it is able to seamlessly handle both missing and categorical data, which is not the case with the implementations of the other benchmarking algorithms. Both the missing data and categorical (mixed) data challenges often faced in the LA datasets, as well as the need to generate highly complex decision boundaries, do seem to indicate that this algorithm should form a part of the suite of algorithms that practitioners consider using in this problem domain.

7. Conclusions

The large volumes of LMS data related to teaching and learning interactions hold hidden knowledge about student learning behaviours. Educational data mining methods have the potential to extract behavioural patterns for the purpose of improving student academic performances via prediction models. Early prediction of students' academic performance enables the identification of potential at-risk students, which provides opportunities for timely intervention to support them while at the same time encouraging those students who are not at risk to get more out of the course via data-driven recommendations and suggestions.

This study considered the problem of developing predictive models that have the capability to operate accurately across disparate courses in identifying at-risk students. Our study demonstrates how such generic and course-agnostic models can be developed in order to overcome the limitations of building multiple models, where each model is tightly coupled with the specific attributes of different courses. The portability of one model across multiple courses is useful because such generic models are less resource-intensive, easier to maintain, and less likely to overfit under certain conditions.

We formulated the machine learning problem as a binary-classification problem that labelled each student as either at risk of failing the course or otherwise. We demonstrated how features can be engineered that are not tightly coupled to the specifics of each course, and thus retain the property of being portable across all course types. Our experiments used Moodle log data, student demographic information and assignment scores from various semesters. Diverse courses were considered in order to robustly evaluate the degree to which our models were generic and portable.

The experiment was carried out using the commonly used Random Forest, Naïve Bayes, Logistic Regression and k-Nearest Neighbours algorithms, as well as the recently developed CatBoost algorithm. Across several performance metrics used in the experiments, the results indicated that the best performance was achieved using CatBoost on our datasets. CatBoost has capabilities in handling categorical features and missing data, while maintaining competitive generalisation abilities compared to current state-of-the-art algorithms. We performed a series of experiments in which we simulated the progression of a teaching semester, considering how early on within a given semester we could accurately identify an at-risk student. For this, our experiments considered data at regular intervals (i.e., at the end of week two, four, six and eight). Our results showed that from as early

as two weeks into a course, our generic, course-agnostic model built using CatBoost was viable for identification of at-risk students and had the potential to reduce academic failure rates through early interventions.

Further, we observed that attributes related to assignment grades (current and prior grades) have a greater impact on model performance relative to other features. Also, students' pre-enrolment data such as study mode, highest school qualification and English proficiency contributed positively to the model prediction. However, one of the most significant challenges in the space of predictive LA is to address how developed models can be effectively deployed across new and diverse courses. Our future work will expand the capabilities of our proposed model to an even broader set of courses.

Author Contributions: Conceptualisation, G.R., T.S. and A.M.; Writing—original draft, G.R., T.S. and A.M.; Writing—review & editing, G.R., T.S. and A.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Conflicts of Interest: The authors declare no conflict of interests.

References

1. The New Media CoNsortium. Available online: <http://www.hp.com> (accessed on 15 June 2021).
2. Junco, R.; Clem, C. Predicting course outcomes with digital textbook usage data. *Internet High. Educ.* **2015**, *27*, 54–63. [[CrossRef](#)]
3. Schumacher, C.; Ifenthaler, D. Features students really expect from learning analytics. *Comput. Hum. Behav.* **2018**, *78*, 397–407. [[CrossRef](#)]
4. Yang, C.C.Y.; Chen, I.Y.L.; Ogata, H. International Forum of Educational Technology & Society Toward Precision Education. *Educ. Technol. Soc.* **2021**, *24*, 152–163. [[CrossRef](#)]
5. Cavus, N. Distance Learning and Learning Management Systems. *Procedia-Soc. Behav. Sci.* **2015**, *191*, 872–877. [[CrossRef](#)]
6. Romero, C.; Ventura, S. Educational Data Mining: A Review of the State of the Art. *IEEE Trans. Syst. Man Cybern. Part C (Appl. Rev.)* **2010**, *40*, 601–618. [[CrossRef](#)]
7. Conijn, R.; Snijders, C.; Kleingeld, A.; Matzat, U. Predicting Student Performance from LMS Data: A Comparison of 17 Blended Courses Using Moodle LMS. *IEEE Trans. Learn. Technol.* **2017**, *10*, 17–29. [[CrossRef](#)]
8. Lust, G.; Elen, J.; Clarebout, G. Students' tool-use within a web enhanced course: Explanatory mechanisms of students' tool-use pattern. *Comput. Hum. Behav.* **2013**, *29*, 2013–2021. [[CrossRef](#)]
9. López-Zambrano, J.; Lara, J.A.; Romero, C. Towards Portability of Models for Predicting Students' Final Performance in University Courses Starting from Moodle Logs. *Appl. Sci.* **2020**, *10*, 354. [[CrossRef](#)]
10. Namoun, A.; Alshantqiti, A. Predicting student performance using data mining and learning analytics techniques: A systematic literature review. *Appl. Sci.* **2021**, *11*, 237. [[CrossRef](#)]
11. Chen, F.; Cui, Y. Utilizing Student Time Series Behaviour in Learning Management Systems for Early Prediction of Course Performance. *J. Learn. Anal.* **2020**, *7*, 1–17. [[CrossRef](#)]
12. Nakayama, M.; Mutsuura, K.; Yamamoto, H. The possibility of predicting learning performance using features of note taking activities and instructions in a blended learning environment. *Int. J. Educ. Technol. High. Educ.* **2017**, *14*, 6. [[CrossRef](#)]
13. Gašević, D.; Dawson, S.; Rogers, T.; Gasevic, D. Learning analytics should not promote one size fits all: The effects of instructional conditions in predicting academic success. *Internet High. Educ.* **2016**, *28*, 68–84. [[CrossRef](#)]
14. Riestra-González, M.; Paule-Ruiz, M.d.; Ortin, F. Massive LMS log data analysis for the early prediction of course-agnostic student performance. *Comput. Educ.* **2021**, *163*, 104108. [[CrossRef](#)]
15. Queiroga, E.; Lopes, J.L.; Kappel, K.; Aguiar, M.S.; Araujo, R.M.; Munoz, R.; Villarroel, R.; Cechinel, C. A Learning Analytics Approach to Identify Students at Risk of Dropout: A Case Study with a Technical Distance Education Course. *Appl. Sci.* **2020**, *10*, 3998. [[CrossRef](#)]
16. Zhao, Q.; Wang, J.-L.; Pao, T.-L.; Wang, L.-Y. Modified Fuzzy Rule-Based Classification System for Early Warning of Student Learning. *J. Educ. Technol. Syst.* **2020**, *48*, 385–406. [[CrossRef](#)]
17. Ramaswami, G.S.; Susnjak, T.; Mathrani, A.; Umer, R. Predicting Students Final Academic Performance using Feature Selection Approaches. In Proceedings of the 2020 IEEE Asia-Pacific Conference on Computer Science and Data Engineering (CSDE), Gold Coast, Australia, 16–18 December 2020. [[CrossRef](#)]
18. Howard, E.; Meehan, M.; Parnell, A. Contrasting prediction methods for early warning systems at undergraduate level. *Internet High. Educ.* **2018**, *37*, 66–75. [[CrossRef](#)]
19. Wolpert, D.H.; Macready, W.G. No Free Lunch Theorems for Optimization. *IEEE Trans. Evol. Comput.* **1997**, *1*, 67–82. [[CrossRef](#)]

20. Tayebinik, M.; Puteh, M. Blended Learning or E-learning? Available online: <http://ssrn.com/abstract=2282881> (accessed on 30 September 2021).
21. Estacio, R.R.; Raga, R.C., Jr. Analyzing students online learning behavior in blended courses using Moodle. *Asian Assoc. Open Univ. J.* **2017**, *12*, 52–68. [[CrossRef](#)]
22. Japkowicz, N.; Stephen, S. The class imbalance problem: A systematic study. *Intell. Data Anal.* **2002**, *6*, 429–449. [[CrossRef](#)]
23. Dorogush, A.V.; Ershov, V.; Gulin, A.; CatBoost: Gradient Boosting with Categorical Features Support. October 2018. Available online: <http://arxiv.org/abs/1810.11363> (accessed on 30 June 2021).
24. Hancock, J.T.; Khoshgoftaar, T.M. CatBoost for big data: An interdisciplinary review. *J. Big Data* **2020**, *7*, 94. [[CrossRef](#)]
25. Mingyu, Z.; Sutong, W.; Yanzhang, W.; Dujuan, W. An interpretable prediction method for university student academic crisis warning. *Complex Intell. Syst.* **2021**, 1–14. [[CrossRef](#)]
26. Breiman, L. Random forests. *Mach. Learn.* **2001**, *45*, 5–32. [[CrossRef](#)]
27. Domingos, P.; Pazzani, M. On the Optimality of the Simple Bayesian Classifier under Zero-One Loss. *Mach. Learn.* **1997**, *29*, 103–130. [[CrossRef](#)]
28. Hosmer, D.W.; Lemeshow, S. *Applied Logistic Regression*; John Wiley & Sons: New York, NY, USA, 2000.
29. Hechenbichler, K.; Schliep, K. Weighted k-Nearest-Neighbor Techniques and Ordinal Classification Projektpartner Weighted k-Nearest-Neighbor Techniques and Ordinal Classification. 2004. Available online: <http://epub.ub.uni-muenchen.de/> (accessed on 4 October 2021).
30. Fabianpedregosa, F.P. Scikit-Learn: Machine Learning in Python Gaël Varoquaux Bertrand Thirion Vincent Dubourg Alexandre Passos Pedregosa, Varoquaux, Gramfort et al. Matthieu Perrot. 2011. Available online: <http://scikit-learn.sourceforge.net> (accessed on 7 October 2021).
31. Ferri, C.; Hernández-Orallo, J.; Modroiu, R. An experimental comparison of performance measures for classification. *Pattern Recognit. Lett.* **2009**, *30*, 27–38. [[CrossRef](#)]
32. Fawcett, T. An introduction to ROC analysis. *Pattern Recognit. Lett.* **2006**, *27*, 861–874. [[CrossRef](#)]
33. Rice, M.E.; Harris, G.T. Comparing effect sizes in follow-up studies: ROC Area, Cohen's d, and r. *Law Hum. Behav.* **2005**, *29*, 615–620. [[CrossRef](#)]
34. Lundberg, S.M.; Allen, P.G.; Lee, S.-I. A Unified Approach to Interpreting Model Predictions. Available online: <https://github.com/slundberg/shap> (accessed on 12 October 2021).

Chapter 6

Supporting Students' Academic

Performance Using Explainable Machine

Learning with Automated Prescriptive

Analytics

The contents of this chapter are from the following article. In accordance with the Multidisciplinary Digital Publishing Institute (MDPI) open access policy, any full text that has been included, is the unmodified accepted article.

@2022 MDPI. Reprinted, with permission, from: Ramaswami, G., Susnjak, T., & Mathrani, A. (2022). Supporting Students' Academic Performance Using Explainable Machine Learning with Automated Prescriptive Analytics. *Big Data and Cognitive Computing*, 6(4), pp. 1-14. <https://doi.org/10.3390/bdcc6040105>

The MDPI open access policy, permits the use, sharing, adaptation, distribution, and reproduction in any medium or format as long as appropriate credit goes to the authors. MDPI does not endorse any of Massey University's products or services. Internal or personal use of this material is permitted. If interested in reprinting/republishing MDPI copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribution, please go to: <https://www.mdpi.com/authors/rights> copyright to learn how to obtain a license from the correct link.

STATEMENT OF CONTRIBUTION DOCTORATE WITH PUBLICATIONS/MANUSCRIPTS

We, the candidate and the candidate's Primary Supervisor, certify that all co-authors have consented to their work being included in the thesis and they have accepted the candidate's contribution as indicated below in the *Statement of Originality*.

Name of candidate:	Gomathy Suganya Ramaswami
Name/title of Primary Supervisor:	Dr. Teo Susnjak
In which chapter is the manuscript /published work:	Chapter 6
Please select one of the following three options:	
<input checked="" type="radio"/> The manuscript/published work is published or in press <ul style="list-style-type: none"> • Please provide the full reference of the Research Output: Ramaswami, G., Susnjak, T., & Mathrani, A. (2022). Supporting Students' Academic Performance Using Explainable Machine Learning with Automated Prescriptive Analytics. <i>Big Data and Cognitive Computing</i>, 6(4), 105. https://doi.org/10.3390/bdcc6040105 	
<input type="radio"/> The manuscript is currently under review for publication – please indicate: <ul style="list-style-type: none"> • The name of the journal: • The percentage of the manuscript/published work that was contributed by the candidate: • Describe the contribution that the candidate has made to the manuscript/published work: 	
<input type="radio"/> It is intended that the manuscript will be published, but it has not yet been submitted to a journal	
Candidate's Signature:	Gomathy Ramaswami <small>Digitally signed by Gomathy Ramaswami Date: 2022.11.07 07:42:59 +13'00'</small>
Date:	07-Nov-2022
Primary Supervisor's Signature:	Teo Susnjak <small>Digitally signed by Teo Susnjak Date: 2022.11.07 15:30:10 +13'00'</small>
Date:	7-Nov-2022

This form should appear at the end of each thesis chapter/section/appendix submitted as a manuscript/publication or collected as an appendix at the end of the thesis.



Article

Supporting Students' Academic Performance Using Explainable Machine Learning with Automated Prescriptive Analytics

Gomathy Ramaswami *, Teo Susnjak and Anuradha Mathrani

School of Mathematical and Computational Sciences, Massey University, Auckland 0632, New Zealand

* Correspondence: g.ramaswami@massey.ac.nz

Abstract: Learning Analytics (LA) refers to the use of students' interaction data within educational environments for enhancing teaching and learning environments. To date, the major focus in LA has been on descriptive and predictive analytics. Nevertheless, prescriptive analytics is now seen as a future area of development. Prescriptive analytics is the next step towards increasing LA maturity, leading to proactive decision-making for improving students' performance. This aims to provide data-driven suggestions to students who are at risk of non-completions or other sub-optimal outcomes. These suggestions are based on *what-if* modeling, which leverages machine learning to model what the minimal changes to the students' behavioral and performance patterns would be required to realize a more desirable outcome. The results of the *what-if* modeling lead to precise suggestions that can be converted into evidence-based advice to students. All existing studies in the educational domain have, until now, predicted students' performance and have not undertaken further steps that either explain the predictive decisions or explore the generation of prescriptive modeling. Our proposed method extends much of the work performed in this field to date. Firstly, we demonstrate the use of model explainability using anchors to provide reasons and reasoning behind predictive models to enable the transparency of predictive models. Secondly, we show how prescriptive analytics based on *what-if* counterfactuals can be used to automate student feedback through prescriptive analytics.

Keywords: machine learning; anchors; counterfactuals; explainable machine learning



Citation: Ramaswami, G.; Susnjak, T.; Mathrani, A. Supporting Students' Academic Performance Using Explainable Machine Learning with Automated Prescriptive Analytics. *Big Data Cogn. Comput.* **2022**, *6*, 105. <https://doi.org/10.3390/bdcc6040105>

Academic Editor: Sabri Pillana

Received: 8 August 2022

Accepted: 25 September 2022

Published: 30 September 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Higher education institutions nowadays are driven by a complex collection of information and educational technologies. This has brought about a new phenomenon: the datafication of learning. Each data point, once aggregated and analyzed, may hold the potential to discover impactful new insights into students' academic profiles and learning behaviors [1]. Learning Analytics (LA) involves the use of a wide selection of institutional data and analytic techniques, ranging from mere descriptive, to predictive, and more recently, prescriptive. Various stakeholders, such as administrators, teaching staff, and students, are now increasingly enabled to act and respond to outputs from data-driven analytics.

Descriptive analytics is the simplest form of data analysis and involves the process of using current and historical data to identify past behaviors and trends from students' data. On the other hand, predictive analytics leverages Machine Learning (ML) algorithms to analyze students' data and generate models that can make forecasts about the likelihood of some phenomena occurring, or inform stakeholders about consequential learning behaviors in the context of LA. While immensely useful, the shortcoming of many of these algorithms is that they produce black-box models which do not offer transparency or insight into the mechanics of these predictions, nor do they provide human-centric explanations of their outputs [2]. This downside often results in distrust in such technologies and a considerable body of work is currently devoted to developing tools that infuse predictive models with explainability of outputs [3]. Meanwhile, prescriptive analytics adds a further layer of

sophistication and is arguably the only component of the analytics suite capable of offering actionable insights by modeling *what-if* scenarios, or counterfactuals. In the context of LA, it is the data-driven evidence generated from counterfactuals that provide automated recommendations to students (and to student advisers) about which adjustments to learning behaviors are most likely to result in improved learning outcomes.

Recent research studies have indicated that effective feedback with specific recommendations on the possible courses of action can lead to self-regulated learning behaviors [4–7]. However, current research in predictive LA has largely been devoid of transparency which can convey to the users what data inputs drive the models towards their conclusions and, crucially, how a model has reasoned in producing the output for a given user. Furthermore, existing research is sparse on automated prescriptive feedback that draws upon data-driven methods to provide advisable behavioral adjustments to students, which may result in more positive outcomes. Meeting this gap has been recognized as an evidence-based pathway for triggering student reflections for maximizing their learning outcomes and ensuring course completions [6].

Taking these two gaps into account, the motivation of this paper is to demonstrate how the potential of predictive analytics can be maximized by communicating the reasoning of how the model arrived at its conclusions. In doing so, we show how prescriptive analytics can be integrated to generate data-driven advice that students can relate to; thus, enabling effectual learning changes that result in increasing the probability of realizing successful outcomes. This integration of different analytics paradigms is the first of its kind to be embedded into a Learning Analytics Dashboard (LAD) currently being piloted at a higher education institution. This study's main contribution is the utilization of approaches that encompass a more complete spectrum of analytics technologies for the LA domain, together with its implementation and demonstration.

This study poses the following research questions:

- How can explainable ML approaches be leveraged to provide effective and human-understandable reasoning behind their conclusions about a student's academic performance?
- How can automatic data-driven feedback and actionable recommendations be derived?

The remainder of this paper consists of five sections. Related work provides an overview of the existing approaches, followed by an analytics workflow that outlines the steps used for processing data in this study's context. In methods, the techniques used for building predictive and prescriptive models are presented and subsequently described in the results section. The discussion provides perspectives on the findings and directly addresses the research questions, followed by a brief summary of the paper's contributions in the conclusion section.

2. Related Work

This section presents recent work undertaken in Educational Data Mining (EDM) research streams for the purposes of predicting student academic performance and for the identification of factors that affect their performance. An overview of emerging ML approaches with their limitations and drawbacks is elaborated upon next.

2.1. Explainable Artificial Intelligence (XAI)

XAI research in the context of ML aims to look inside black-box predictive models for extracting information or explanations as to why a model has come to a given conclusion. In addition to providing tools that can build more trust and accountability, XAI assists with debugging and identifying bias in ML.

Predicting students' learning behaviors and their final outcomes is regarded as one of the most important tasks in the EDM field. Prior studies have proposed several data mining approaches to predict academic performance, which would then be followed up by notifying the instructor about who the at-risk students are, which in turn opens opportunities for the instructor to intervene and provide the student with learning pathways for

improving their performance. It is common for students' previous academic performance, demographics, and LMS usage-related features to be used for this predictive task; however, all previous studies' focus has been merely on either conveying to students their at-risk status of non-completion in a course [8–11] or in delivering the probability of attaining specific grades [12].

None of these predictive LA applications adequately explain to users how the models work, or how their predictions were generated; rather, they simply provide the predicted outcomes as feedback to users. Although such feedback might be helpful to some extent, it does not provide meaningful personalized insights or actionable information about the reasons behind the predictions. This lack of interpretability and explainability of the models makes it difficult for end-users to trust the system with which they are interacting. Moreover, when a model produces unexpected or erroneous output, the lack of trust often results in increased skepticism and possibly even a rejection on the part of the end-user [6]. These errors may have negative side effects as well, such as in instances when certain actions or decisions affecting others are taken, which might be based on false premises arising from misclassifications [13]. Thus, it has become a *sine qua non* to investigate how the inference mechanism or the decisions of ML models can be made transparent to humans so that Artificial Intelligence (AI) can become more acceptable to users from different application domains [14]. This has led to the establishment of a relatively new research field, namely *eXplainable AI*. The XAI domain researches and develops tools and techniques that enable the interpretability of autonomous decision-making systems through outputs in the form of simplified textual and visual representations that can be understood by human users [8].

2.2. Post Hoc Explanations of Machine Learning Models

Extracting explanations of ML models after they are induced is crucial for users (or students in this context) to apprehend so that they can act on algorithmic predictions in an informed manner. Moreover, explanations of an ML model's predictions have found many uses, including understanding which features are most important, debugging the models, and increasing trust through transparency [5]. In addition to furthering the societal acceptance of the recommendations that are based upon algorithmic decisions, they also provide alternate suggestions for overcoming unfavorable predictions.

Meanwhile, when it comes to prescriptive analytics, several studies have leveraged different forms of messaging and recommendation techniques to enhance their prescriptive capabilities (e.g., [8,9,14]). However, none of the existing studies in EDM employed prescriptive modeling algorithms for generating automated personalized recommendations to inform students of specific behavioral changes [6]. In a recent study [6], we implemented an interpretability component and operationalized a LA dashboard that provided post hoc explanations of the prescribed recommendations that were tailored for individual students. We believe this to be the first published state-of-the-art LA dashboard to have incorporated both automated prescriptive analytics as well as transparent predictive models in a student learning context.

3. Analytics Workflow

This section presents the conventional steps involved in the predictive analytics process flow, to which we add two further steps which introduce both model transparency and prescriptive analytics.

3.1. Standard Predictive Analytics Workflow

The general workflow for predictive analytics for generating ML models involves acquiring relevant datasets as the first step (refer to Figure 1). Next, the pre-processing of the raw (unclean) data is performed to prepare this data, resulting in reliable datasets to enable subsequent analyses. In this step, usually, the total number of features (or variables) is reduced and only the most relevant or complete features are retained. In the data

transformation step, additional features are also engineered as derivatives of the raw data to produce richer descriptors that are more correlated with the target (dependent) variable.

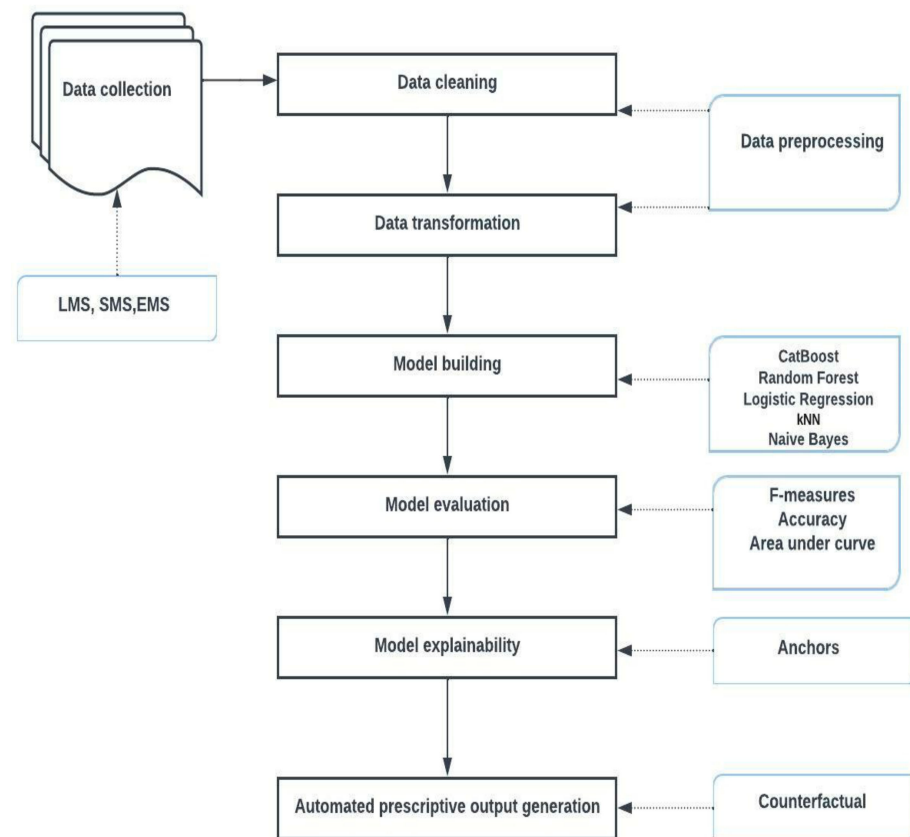


Figure 1. Analytics workflow followed in this study.

Predictive models are subsequently induced based on the selected features in conjunction with a suite of ML algorithm implementations. The resulting predictive models are then evaluated using various evaluation measures for their generalizability. In the case of this study, a range of metrics was used, and finally, the best predictive models were selected based on these metrics.

3.2. Model Explainability

Once the predictive model has been generated, we propose the introduction of a new step in the workflow, which generates an explanation of how a given model has arrived at a particular prediction for a specific student. When using black-box models, there are technologies available that generate *proxy* models which emulate the behavior of the underlying models, since the purpose is to simplify and present the overall model behavior in human-understandable terms [15]. Technologies such as Anchors [16], SHAP [17], and LIME [18] can typically be leveraged to generate these simplified explanations.

Anchors are model-agnostic, i.e., they can be applied to any prediction model without requiring knowledge about the internals of the original model. Anchors uses a perturbation-based strategy to generate easy-to-understand explanations for predictions by emulating the behavior of the original, which are then expressed as IF-THEN rules. In order for them to be meaningful, the Anchors need to have a high precision and coverage score. Precision is the proportion of data points in the region defined by the anchor that have the same class as the data point being explained, while coverage describes how many data points an anchor's decision rule applies to [19].

3.3. Prescriptive Analytics

The final proposed step involves generating data-driven prescriptive feedback. In this step, the desired target outcome is selected for a given student (i.e., a successful outcome for a student whose prediction is otherwise), and an ML tool is applied which computes the smallest set of changes that would need to take place amongst the selected features to achieve the desired predicted output [20]. Counterfactuals are chosen for generating the automated prescriptive output in this study.

4. Methods

4.1. Dataset and Features

The datasets used in this study were extracted from courses offered at an Australasian higher education institution involving a blended learning environment. These comprised data describing learning activities from Moodle (i.e., the learning management system (LMS)), various assessment grades from the Student Management System (SMS), and the demographic and pre-academic data from the Enrolment Management System (EMS). The raw data from the LMS log files were processed to engineer features used for building various analytics components within the predictive and prescriptive models.

During this process, individual student behavioral attributes describing learning patterns were extracted from the raw data and converted into a format that relativized each student’s value with respect to the average patterns of the student’s cohort. Z-score standardization was used to achieve this. This transformation technique had the effect of producing features that were generic and thus applicable across different courses and cohorts. The final features used for model development are shown in Table 1. These consist of assignment score, assignment deviation score, LMS engagement score, LMS engagement deviation score, prior course scores, resources viewed, citizenship, and demographic information such as age and the gender of the students.

Table 1. Feature descriptions.

Feature Name	Description
Assignment score	Assignment score received by a student
Assignment deviation score	The Z-score of the student’s mean assignment score as a deviation from the cohort mean
LMS engagement score	The count of all activities performed by a student on the Moodle platform.
LMS deviation score	The engagement score expressed as a Z-score of a student as a deviation from the cohort mean
Prior course scores	The mean score achieved by a student from across all previous course scores
Resources viewed	Resources viewed by a student on the Moodle
Citizenship	The nationality of the student
Age	Age of a person
Gender	Gender of the student
English proficiency test	English proficiency
Study mode	Study by distance/online or on-campus

For the model training, we defined two categories of students, namely low-risk and high-risk. The threshold used for assigning these labels to the students depended on the students’ final course mark. Students with a course mark of 60% or less were labeled as high-risk, while the remainder were considered low risk. These categories alluded to the risk level of failing a course.

4.2. Tools

We used a mixture of Python’s scikit-learn library [21] and the separate CatBoost [22] library implementation for classification algorithms covering the predictive analytics component. For the explainability component of our analytics workflow, we used Anchors

proxy models, and specifically the Python [21] Anchors library implementation. An anchor explanation is simply a translation of the black box into a set of rule-based outputs based on *if-else* conditions and is thus intuitive, and easy to comprehend. In order to describe high-level model mechanics, we used the Shapely Additive Explanations (SHAP). SHAP calculates the local feature importance for every observation. The SHAP method constructs an additive interpretation model based on the Shapley value. The Shapley value measures the marginal contribution of each feature to the entire cooperation.

For generating prescriptive analytics outputs leveraging counterfactual explanations, we used Python's DiCE (Diverse Counterfactual Explanation) [20] implementation in our experiments. DiCE computes a set of diverse counterfactual explanations by finding candidate feature vectors that are close to the query instance but with an opposite prediction [23].

4.3. Machine Learning Algorithms and Evaluation

We employed five ML algorithms for the prediction of learning outcomes (low-risk/high-risk), namely Logistic Regressions (LR) [24], k-Nearest Neighbors (KNN) [25], Random Forest (RF) [26], Naïve Bayes (NB) [27], and CatBoost. The purpose was to compare and evaluate these classifiers thoroughly to determine which one consistently performed better on this dataset. To ensure that the final model (that has been described in detail in our published work [28]) generalizes well, the predictive model's performance was evaluated using accuracy, precision, recall, F-measure, and area under the curve (AUC).

A modified k-fold cross-validation approach was used to evaluate the models. Our dataset consisted of seven courses and a total of 10 separate deliveries of those courses. The models were trained on nine course deliveries and tested against the remaining hold-out course offering. The process was repeated 10 times with a different combination of training and hold-out courses in order to arrive at our final, aggregated evaluation scores for our models. We can observe from the results (Table 2) that the overall accuracy of all the algorithms has ranged between 65% and 75%, with the result from CatBoost clearly outperforming the other algorithms on these datasets. The accuracies of the bottom three algorithms, namely Naïve Bayes, Random Forest, and Logistic Regression, did not exhibit significantly divergent accuracy results.

Table 2. Performance scores of various classifiers.

Classifiers	F-Measure	Accuracy	AUC
CatBoost	0.77 ± 0.024	75 ± 2.1	0.87 ± 0.023
Random Forest	0.67 ± 0.025	67 ± 2.4	0.74 ± 0.015
Naïve Bayes	0.67 ± 0.023	68 ± 2.3	0.71 ± 0.034
Logistic Regression	0.68 ± 0.031	67 ± 3	0.73 ± 0.025
k-Nearest Neighbors	0.71 ± 0.02	71 ± 2.4	0.72 ± 0.022

Figure 2 depicts the F-measure for various classifiers on the hold-out dataset, depicting the scores at different snapshots in time, where the snapshots were defined by the week number in a given semester at which point the predictions were made. The snapshots were 2, 4, 6, and 8-week time points in a semester. The figure shows that, on average, the predictive accuracy of the models improved as the semester progressed, and more data about the students was gathered. The final accuracy from all the test datasets was displayed to the students as a measure of confidence in the reliability of the underlying model.

4.4. Feature Importance

Feature importance analysis is an important concept of machine learning as it helps to estimate which features are making the most impact on a model's decision-making. Figure 3 shows the feature importance of our model, listing the most important features at the top, with the size of the bars indicating the magnitude of impact. The figure indicates that the top three features in the model's reasoning all heavily relied on students' academic performance in a prior course.

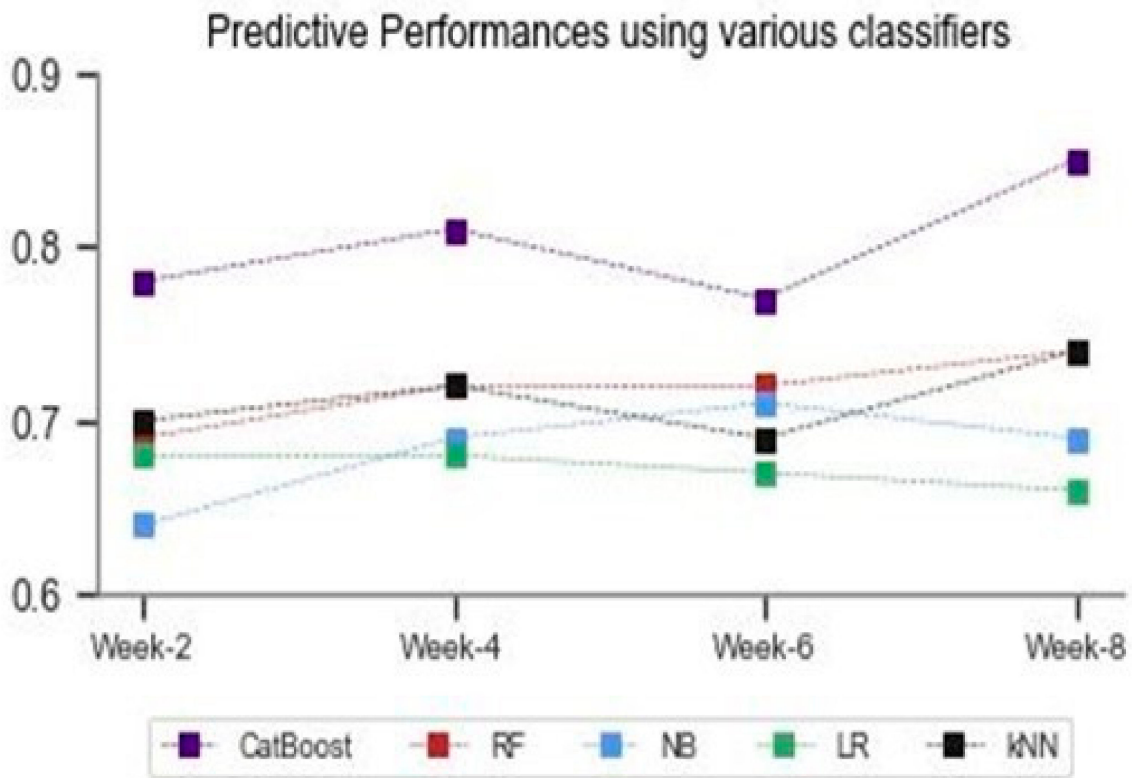


Figure 2. Predictive performances of various classifiers.

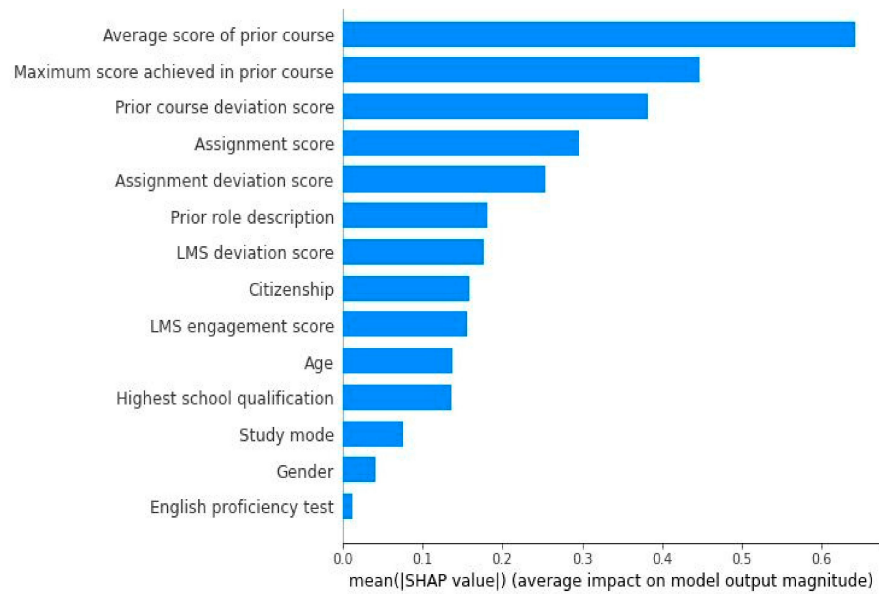


Figure 3. Feature importance plot.

In order to extract some additional insights into the mechanics of the models, we use feature dependence plots in Figure 4. A dependence scatter plot explores the interaction effects between different features on the predicted outcome of a model. The x -axis denotes the value of a target feature, and the y -axis is the corresponding SHAP value for that feature, which relates directly to the effect it has on the final prediction.

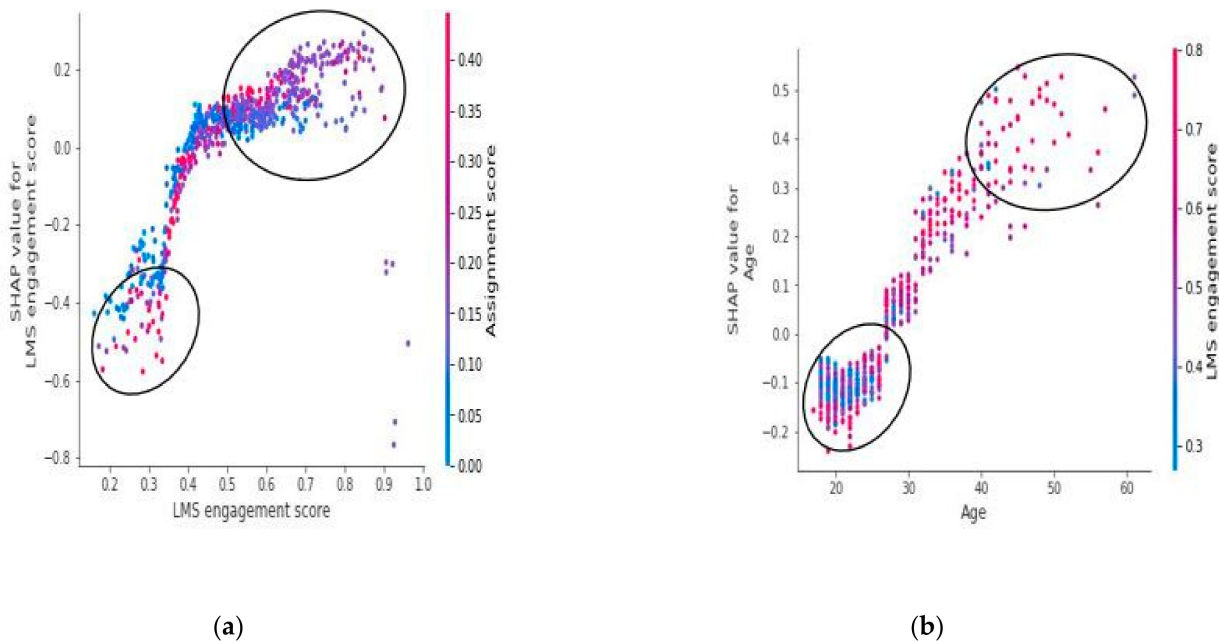


Figure 4. Two SHAP scatter dependence plots. (a) interaction between LMS engagement score and assignment score. (b) interaction between age and LMS engagement score.

Figure 4a indicates that there is a positive correlation between LMS engagement scores and assignment scores. Two noteworthy patterns emerge from Figure 4a. First, we can see that students who score highly on the assignment scores but exhibit poor engagement levels with the LMS receive strongly negative outcome predictions by the model. We also see there is a threshold of 0.4 for the LMS engagement score, and those that score above this threshold are positively correlated with the model's predictions for successful outcomes, which is amplified further for those with higher assignment scores.

Figure 4b denotes feature interaction between the LMS engagement score and the age of the student. It can be seen from the plot as the age of a student increases, the effect on the model predictions for positive outcomes becomes stronger. From approximately age 26 onwards, increases in the student age carry a stronger positive effect on model predictions for positive outcomes until age 40, from which point there do not appear to be any further increasing positive effects. The students who are most at-risk are those in their early twenties, with LMS engagement scores having no clear positive effect on prediction outcomes for this student demographic.

5. Result

This section presents the results of transforming a black-box model prediction into explainable outputs, and it also shows how prescriptive feedback can be generated from *what-if* analyses. The results are a demonstration taken from two hypothetical students (Students A and B) who were identified as high-risk by the predictive model that was trained as outlined in Figure 1 and described in the methodology section.

5.1. Model Explainability Example Using Anchors

Figure 5 shows how the model has reasoned in concluding that Student A and Student B as being at high risk of course non-completion. The output is generated by the Anchors method, which creates an intermediate (proxy) model that emulates the behavior of the actual underlying black-box model and converts it into a human-understandable sequence of predicates. The rule-based classifier can be understood as categorizing a student as high-risk if every predicate evaluates as true. The conversion process from a black box into an intermediate proxy model not only provides insights into exact feature value thresholds that contribute to the eventual predicted outcome but also demonstrates which particular

features are important in the classification process. In this example for Student A (Figure 5a), the features assignment scores, and the online engagement with the Moodle online learning environment have influenced the model to classify the student into the high-risk category. The simplified output of the proxy model is able to clearly communicate that both the low assignment scores and the student’s lower engagement score compared with the student’s cohort, as being the key drivers of the prediction while also showing the exact thresholds in those values.

<p>Prediction explainability rule</p> <p>Assignment scores \leq 63% AND</p> <p>Assignment deviation scores \leq 3.17 AND</p> <p>LMS engagement deviation score \leq -0.49 AND</p> <p>LMS engagement score \leq 0.42</p>	<p>Prediction</p> <p>High-risk</p>
--	---

(a)

<p>Prediction explainability rule</p> <p>Assignment scores \leq 47% AND</p> <p>Prior course scores \leq 60% AND</p> <p>LMS engagement deviation score \leq 0.20 AND</p> <p>LMS engagement score \leq 0.27</p>	<p>Prediction</p> <p>High-risk</p>
--	---

(b)

Figure 5. (a) Anchor explanation for Student A. (b) Anchor explanation for Student B.

However, in the second example for Student B in Figure 5b, we can see that the model explanation deviates from that of the previous example; in contrast, it draws upon two different features in order to explain the high-risk classification for this student. The explanation utility, in this case, uses the student’s prior course scores and the online resources viewed features to arrive at the student’s eventual classification. Herein lies the potential of prescriptive analytics that is able to leverage different combinations of business rules in order to provide tailored prescriptive feedback to the students. Students can thus be better advised on what remedial actions could be taken, which might result in different and more positive outcomes.

5.2. Prescriptive Analytics Example Using Counterfactuals

The proxy models in the previous section are particularly useful in offering explanations of their reasoning to relevant stakeholders. Some prescriptive insights can be extracted from them; however, in this section, we demonstrate how with the aid of an additional technique, we can generate more precise and specific prescriptive feedback. To demonstrate this, we use the same hypothetical students as in the previous section as examples. The

outputs in Table 3 are generated using counterfactuals that model what-if scenarios by calculating what minimal changes in the existing student’s feature values need to occur in order for an opposite outcome (low-risk in our context) to be predicted [29].

Table 3. (a) Counterfactual explanation for Student A. (b) Counterfactual explanation for Student B.

(a)						
	Query Instance (Original Outcome: High-Risk)		Counterfactual Set (New Outcome: Low-Risk)			
	Activity	Value	Activity	Value		
Student A	Assignment score	39%	Assignment score	61%		
	Assignment deviation score	− 1.1	Assignment deviation score	− 1.1		
	LMS engagement score	0.22	LMS engagement score	0.52		
	LMS engagement deviation score	− 0.49	LMS engagement score	− 0.49		
(b)						
	Query Instance (Original Outcome: High-Risk)		Counterfactual-Set1 (New Outcome: Low-Risk)		Counterfactual-Set2 (New Outcome: Low-Risk)	
	Activity	Value	Activity	Value	Activity	Value
Student B	Assignment score	31%	Assignment score	63%	Assignment score	63%
	Prior course score	55%	Prior course score	60%	Prior course score	55%
	Resources viewed score	0.12	Resources viewed score	0.32	Resources viewed score	0.42
	LMS engagement score	0.22	LMS engagement score	0.22	LMS engagement score	0.35

Your current assignment average is 39% and on average your engagement into Stream 3 times a week. To increase your likelihood of succeeding in the course, we recommend that you strive towards achieving at least 61% in your next assignment submission and to increase your online engagement levels to at least 5 times a week.

(a)

Your current assignment average is 31% and on average your engagement into Stream 2 times a week and the resources viewed was 1 time a week. To increase your likelihood of succeeding in the course, we recommend that you strive towards achieving at least 63% in your next assignment submission and to increase your online engagement levels and resources viewed to at least 4 times a week.

(b)

Figure 6. (a) Counterfactual explanation translation for Student A. (b) Counterfactual explanation translation for Student B.

The tables show Student A and Student B values termed as a *query instance* (together with the original classification outcome) and the generated counterfactual set (desired

outcome for the student, which is low risk in this context). The values column in the tables represents the actual and the modeled feature values with their corresponding feature names. Table 3a shows a single counterfactual set being generated for Student A, while Table 3b demonstrates how multiple counterfactual sets can be generated. All the generated counterfactual sets are based on the underlying black-box model and demonstrate minimal changes to the various feature values that would be required to flip the classification of those students from high to low-risk, which can then directly be used for prescriptive feedback. For Student A (Table 3a), the risk level from “high” to “low” would occur if there is an improvement in scores in the upcoming assignment and online engagement with Moodle; however, for the hypothetical Student B (Table 3b), the first pathway suggests that the change in prior course scores along with assignment scores can help the student to change their risk level; albeit, in this context, a change in prior course scores cannot really be actioned for the current course, though it can indicate to the student how they could be classified as lower risk for their subsequent course if they achieve a certain score in their present course. In the second alternative set, it is suggested that the changes in assignment score, resources viewed, and engagement with Moodle can help the corresponding student to fall into the low-risk level.

Once the above counterfactual explanations have been induced, they can then be translated into a human-readable format as seen below (Figure 6) in order to provide students with clear, precise, and actionable suggestions on what adjustments to their learning behavior and performance are most likely to assist them in realizing positive outcomes. The translation itself can also be automated with the assistance of natural language processing tools with the confidence that the prescriptive feedback given to students is data-driven and reliable, provided that feasible counterfactual sets are selected.

6. Discussion

ML is at the core of many recent advances in science and technology. Whether humans are directly using ML classifiers as tools, or are deploying models within other products, a vital concern remains: if users do not trust a model or a prediction, they will not use it. This is particularly relevant in settings where users are expected to make decisions based on the outputs of machine learning models. Besides helping to debug ML models, explanations of black-box models improve the interpretability and trustworthiness of algorithmic decisions and enhance human decision-making [17]. We propose the use of anchors to identify the influencing factors in the predictive models to explain to students in an easily understandable way how the model works and how the prediction is generated [30].

Not only should the explanatory model provide customized individual suggestions to users, but they should also prioritize which changes in behavior and which learning strategies the student should adopt that will most likely translate into favorable results. This should take the guesswork out of the equation and help students make good decisions. Expanding the LA models to include prescriptive analytics that integrates counterfactuals and automated customized suggestions to students creates an approach that further supports students and helps them to achieve the best possible learning outcomes.

Our study highlighted the importance of moving beyond black-box approaches and the value of openly explaining to students how the prediction of their outcome or risk status was made, and furthermore providing customized and actionable suggestions based on real data to help the students work towards a favorable outcome. If students first understand why they have been classified in a certain way and then provided with clear and precise steps they can take to change a potentially unfavorable outcome, they are more likely to make the necessary changes. This study, as far as we know, is the first in the education domain in the field of LA dashboards to combine these approaches.

7. Conclusions

The prediction of academic performance is considered one of the most popular tasks in the field of LA. However, we found from prior studies that a disproportionate amount

of attention is given to merely creating predictive models, and an insufficient amount of focus is placed on addressing model interpretability and the explainability of the predictive outputs. The latter lowers the utility of this technology, and over time erodes the trust of users. Furthermore, the current focus of machine learning technologies mostly stops at the predictive analytics step and does not provide stakeholders with the translation to actionable insights that can be derived from more sophisticated approaches which yield prescriptive analytics.

Hence the current study not only aimed to demonstrate how model interpretability and explainability of the individual predictions to students can be embedded into LA systems, but crucially, we demonstrate how this can be taken one step further and converted into counterfactuals, from which prescriptive interventions can be provided to the students. We present a demonstration of how LA systems can be developed which encompass a more comprehensive machine learning approach towards providing support to students which integrates both predictive and prescriptive analytics while maximizing transparency and trust in the underlying data-driven technologies.

The developed system presented in this study has been integrated into a learning analytics dashboard, and a pilot study is currently being conducted with the students to analyze the effectiveness of the dashboard with student final outcomes.

Author Contributions: Conceptualization, G.R., T.S. and A.M.; Writing—original draft, G.R., T.S. and A.M.; Writing—review and editing, G.R., T.S. and A.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: This study has received approval from the Institutional Human Ethics Committee to conduct the research.

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: Not applicable.

Conflicts of Interest: The authors declare no conflict of interests.

References

1. Jones, K.M.L. Learning analytics and higher education: A proposed model for establishing informed consent mechanisms to promote student privacy and autonomy. *Int. J. Educ. Technol. High. Educ.* **2019**, *16*, 24. [[CrossRef](#)]
2. Lu, J.; Chen, W.; Ma, Y.; Ke, J.; Li, Z.; Zhang, F.; Maciejewski, R. Recent progress and trends in predictive visual analytics. *Front. Comput. Sci.* **2017**, *11*, 192–207. [[CrossRef](#)]
3. Linardatos, P.; Papastefanopoulos, V.; Kotsiantis, S. Explainable AI: A Review of Machine Learning Interpretability. *Methods Entropy* **2020**, *23*, 18. [[CrossRef](#)] [[PubMed](#)]
4. Lepenioti, K.; Bousdekis, A.; Apostolou, D.; Mentzas, G. Prescriptive analytics: Literature review and research challenges. *Int. J. Inf. Manag.* **2020**, *50*, 27–70. [[CrossRef](#)]
5. Algayres, M.; Triantafyllou, E. Online Environments for Supporting Learning Analytics in the Flipped Classroom: A Scoping Review. In Proceedings of the 18th European Conference on e-Learning, Copenhagen, Denmark, 7–8 November 2019; p. 8. [[CrossRef](#)]
6. Susnjak, T.; Ramaswami, G.S.; Mathrani, A. Learning analytics dashboard: A tool for providing actionable insights to learners. *Int. J. Educ. Technol. High. Educ.* **2022**, *19*, 12. [[CrossRef](#)] [[PubMed](#)]
7. Li, Q.; Li, Z.; Han, J.; Ma, H. Quality Assurance for Performing Arts Education: A Multi-Dimensional Analysis Approach. *Appl. Sci.* **2022**, *12*, 4813. [[CrossRef](#)]
8. Wells, L.; Bednarz, T. Explainable AI and Reinforcement Learning—A Systematic Review of Current Approaches and Trends. *Front. Artif. Intell.* **2021**, *4*, 550030. [[CrossRef](#)] [[PubMed](#)]
9. Fleur, D.S.; van den Bos, W.; Bredeweg, B. Learning Analytics Dashboard for Motivation and Performance. In *International Conference on Intelligent Tutoring Systems*; Springer: Cham, Switzerland, 2020. [[CrossRef](#)]
10. Baneres, D.; Rodriguez, M.E.; Serra, M. An Early Feedback Prediction System for Learners At-Risk Within a First-Year Higher Education Course. *IEEE Trans. Learn. Technol.* **2019**, *12*, 249–263. [[CrossRef](#)]
11. Kokoç, M.; Altun, A. Effects of learner interaction with learning dashboards on academic performance in an e-learning environment. *Behav. Inf. Technol.* **2021**, *40*, 161–175. [[CrossRef](#)]
12. Li, Y.; Gou, J.; Fan, Z. Educational data mining for students' performance based on fuzzy C-means clustering. *J. Eng.* **2019**, *2019*, 8245–8250. [[CrossRef](#)]

13. Valle, N.; Antonenko, P.; Valle, D.; Sommer, M.; Huggins-Manley, A.C.; Dawson, K.; Kim, D.; Baiser, B. Predict or describe? How learning analytics dashboard design influences motivation and statistics anxiety in an online statistics course. *Educ. Technol. Res. Dev.* **2021**, *69*, 1405–1431. [[CrossRef](#)] [[PubMed](#)]
14. Islam, M.R.; Ahmed, M.U.; Barua, S.; Begum, S. A systematic review of explainable artificial intelligence in terms of different application domains and tasks. *Appl. Sci.* **2022**, *12*, 1353. [[CrossRef](#)]
15. Loyola-Gonzalez, O. Black-box vs. white-box: Understanding their advantages and weaknesses from a practical point of view. *IEEE Access* **2019**, *7*, 154096–154113. [[CrossRef](#)]
16. Ribeiro, M.T.; Singh, S.; Guestrin, C. Anchors: High-Precision Model-Agnostic Explanations. In Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, New Orleans, LA, USA, 2–7 February 2018; pp. 1527–1535. Available online: www.aaai.org (accessed on 1 November 2021).
17. Lundberg, S.M.; Allen, P.G.; Lee, S.-I. A Unified Approach to Interpreting Model Predictions. Available online: <https://github.com/slundberg/shap> (accessed on 12 October 2021).
18. Ribeiro, M.T.; Singh, S.; Guestrin, C. Why Should I Trust You? In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 13–17 August 2016; Association for Computing Machinery: New York, NY, USA; pp. 1135–1144. [[CrossRef](#)]
19. Holzinger, A.; Saranti, A.; Molnar, C.; Biecek, P.; Samek, W. Explainable AI methods-abrief overview. In *International Workshop on Extending Explainable AI Beyond Deep Models and Classifiers*; Springer: Cham, Switzerland, 2022; pp. 13–38. [[CrossRef](#)]
20. Mothilal, R.K.; Sharma, A.; Tan, C. Explaining machine learning classifiers through diverse counterfactual explanations. In Proceedings of the FAT* 2020—2020 Conference on Fairness, Accountability, and Transparency, Barcelona, Spain, 27–30 January 2020; pp. 607–617. [[CrossRef](#)]
21. Pedragosa, F.; Varoquaux, G.; Gramfort, A.; Thirion, B.; Dubourg, V.; Passos, A.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Perrot, M.; et al. Scikit-Learn: Machine Learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830. Available online: <http://scikit-learn.sourceforge.net> (accessed on 7 October 2021).
22. Dorogush, A.V.; Ershov, V.; Gulin, A. CatBoost: Gradient Boosting with Categorical Features Support. *arXiv* **2018**, arXiv:1810.11363v1. Available online: <http://arxiv.org/abs/1810.11363> (accessed on 1 November 2021).
23. Bhatt, U.; Xiang, A.; Sharma, S.; Weller, A.; Taly, A.; Jia, Y.; Ghosh, J.; Puri, R.; Moura, J.M.; Eckersley, P. Explainable machine learning in deployment. In Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, Barcelona, Spain, 27–30 January 2020; pp. 648–657. [[CrossRef](#)]
24. Hosmer, D.W.; Lemeshow, S. *Applied Logistic Regression*; John Wiley & Sons: New York, NY, USA, 2000.
25. Hechenbichler, K.; Schliep, K. Weighted k-Nearest-Neighbor Techniques and Ordinal Classification. 2004. Available online: https://epub.uni-muenchen.de/1769/1/paper_399.pdf (accessed on 1 November 2021).
26. Breiman, L. Random forests. *Mach. Learn.* **2001**, *45*, 5–32. [[CrossRef](#)]
27. Domingos, P.; Pazzani, M. On the Optimality of the Simple Bayesian Classifier under Zero-One Loss. *Mach. Learn.* **1997**, *29*, 103–130. [[CrossRef](#)]
28. Ramaswami, G.; Susnjak, T.; Mathrani, A. On Developing Generic Models for Predicting Student Outcomes in Educational Data Mining. *Big Data Cogn. Comput.* **2022**, *6*, 6. [[CrossRef](#)]
29. Wachter, S.; Mittelstadt, B.; Russell, C. Counterfactual Explanations Without Opening the Black Box: Automated Decisions and the GDPR. *SSRN Electron. J.* **2017**, *31*, 841. [[CrossRef](#)]
30. Fernández, R.R.; de Diego, I.M.; Aceña, V.; Fernández-Isabel, A.; Moguerza, J.M. Random forest explainability using counterfactual sets. *Inf. Fusion* **2020**, *63*, 196–207. [[CrossRef](#)]

Chapter 7

Learning analytics dashboard: a tool for
providing actionable insights to learners

The contents of this chapter are from the following article. In accordance with the SpringerLink open access policy, any full text that has been included, is the unmodified accepted article.

@2022 SpringerLink. Reprinted, with permission, from: Susnjak, T., Ramaswami, G. S., & Mathrani, A. (2022). Learning analytics dashboard: a tool for providing actionable insights to learners. *International Journal of Educational Technology in Higher Education*, 19(1), 1-23. <https://doi.org/10.1186/s41239-021-00313-7>.

The SpringerLink open access policy, permits the use, sharing, adaptation, distribution and reproduction in any medium or format as long as appropriate credit goes to the authors. SpringerLink does not endorse any of Massey University's products or services. Internal or personal use of this material is permitted. If interested in reprinting/republishing Springerlink copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribution, please go to:

<https://educationaltechnologyjournal.springeropen.com/submission-guidelines/> copyright to learn how to obtain a license from the correct link.

STATEMENT OF CONTRIBUTION DOCTORATE WITH PUBLICATIONS/MANUSCRIPTS

We, the candidate and the candidate's Primary Supervisor, certify that all co-authors have consented to their work being included in the thesis and they have accepted the candidate's contribution as indicated below in the *Statement of Originality*.

Name of candidate:	Gomathy Suganya Ramaswami
Name/title of Primary Supervisor:	Dr. Teo Susnjak
In which chapter is the manuscript /published work:	Chapter 7
Please select one of the following three options:	
<input checked="" type="radio"/> The manuscript/published work is published or in press <ul style="list-style-type: none"> • Please provide the full reference of the Research Output: Susnjak, T., Ramaswami, G. S., & Mathrani, A. (2022). Learning analytics dashboard: a tool for providing actionable insights to learners. <i>International Journal of Educational Technology in Higher Education</i>, 19(1), 1-23. https://doi.org/10.1186/s41239-021-00313-7 	
<input type="radio"/> The manuscript is currently under review for publication – please indicate: <ul style="list-style-type: none"> • The name of the journal: • The percentage of the manuscript/published work that was contributed by the candidate: • Describe the contribution that the candidate has made to the manuscript/published work: 	
<input type="radio"/> It is intended that the manuscript will be published, but it has not yet been submitted to a journal	
Candidate's Signature:	Gomathy Ramaswami <small>Digitally signed by Gomathy Ramaswami Date: 2022.11.07 07:42:59 +13'00'</small>
Date:	07-Nov-2022
Primary Supervisor's Signature:	Teo Susnjak <small>Digitally signed by Teo Susnjak Date: 2022.11.07 15:30:10 +13'00'</small>
Date:	7-Nov-2022

This form should appear at the end of each thesis chapter/section/appendix submitted as a manuscript/publication or collected as an appendix at the end of the thesis.

RESEARCH ARTICLE

Open Access



Learning analytics dashboard: a tool for providing actionable insights to learners

Teo Susnjak^{*} , Gomathy Suganya Ramaswami and Anuradha Mathrani

^{*}Correspondence:
t.susnjak@massey.ac.nz
School of Mathematical
and Computational Sciences,
Massey University, Auckland,
New Zealand

Abstract

This study investigates current approaches to learning analytics (LA) dashboarding while highlighting challenges faced by education providers in their operationalization. We analyze recent dashboards for their ability to provide actionable insights which promote informed responses by learners in making adjustments to their learning habits. Our study finds that most LA dashboards merely employ surface-level descriptive analytics, while only few go beyond and use predictive analytics. In response to the identified gaps in recently published dashboards, we propose a state-of-the-art dashboard that not only leverages descriptive analytics components, but also integrates machine learning in a way that enables both predictive and prescriptive analytics. We demonstrate how emerging analytics tools can be used in order to enable learners to adequately interpret the predictive model behavior, and more specifically to understand how a predictive model arrives at a given prediction. We highlight how these capabilities build trust and satisfy emerging regulatory requirements surrounding predictive analytics. Additionally, we show how data-driven prescriptive analytics can be deployed within dashboards in order to provide concrete advice to the learners, and thereby increase the likelihood of triggering behavioral changes. Our proposed dashboard is the first of its kind in terms of breadth of analytics that it integrates, and is currently deployed for trials at a higher education institution.

Keywords: Dashboard, Learner analytics, Actionable insights, Model interpretability, Explainable AI, Counterfactuals

Introduction

Analytics technologies have proliferated across many sectors of society for extracting data-driven insights, improving decision making and driving innovation. The tertiary educational sector is particularly in-tune with the advantages that data analytics can offer, and generally, seeks to leverage these advances. Deployment of analytics technologies is becoming increasingly important as this sector is undergoing disruptions across different parts of the world, as well as due to the COVID-19 pandemic crisis (Aristovnik et al., 2020). The current pandemic responses have shifted education delivery to online modes, further accelerating ongoing disruptions. The education sector is already facing financial and competitive pressures (Muhammad et al., 2020) in some regions, and this global shift to online learning has amplified them even more. These shifts have altered

the competitive landscape between universities in countries with low or non-existent subsidies with those located in countries that have strong government support by bringing them into direct competition as geographic and physical boundaries now have a diminished relevance. This has been accentuated by the continuing rise in higher education costs (Blankenberger & Williams, 2020), together with the questioning of the value proposition that higher education offers in many of the available qualifications. These challenges call for a need to adapt and do things differently. For these reasons, data analytics with its use of innovative products has risen to become one of the key enablers in the educational sector. Analytics in this sector focuses on maximizing student retention rates by identifying the at-risk students in the early stages and then by initiating interventions, whilst improving the quality of the educational experience for the learners.

One of the analytics innovation products that have been deployed to benefit the enrolled learners, are Learning Analytics dashboards (LADs) which have been operationalized in numerous institutions (Table 1) in the last few years. The purpose of existing LADs is not dissimilar to the dashboards widely used in industry. They aim to provide learners with a snapshot of how they are progressing in their courses. Graphical displays highlight trends in the learners' academic and engagement levels through the digital footprints generated by learners and provide them with a basis for awareness, reflection and new insights (Yoo & Jin, 2020). In more advanced cases, the dashboards are designed to make predictions on where learners are likely to end up in respect to meeting learning outcomes based on their current trajectory. Analytics tools reveal learner insights that could prompt a reflection process that would otherwise not be realized. It is assumed that these reflections may (in some cases) trigger positive behavioral changes that support learners in maximizing learning outcomes and course completions as well as retention rates.

However, operationalizing these types of analytics products is accompanied with numerous challenges. Given the financial investment and the human resource effort involved in their productionization (Mahroeian et al., 2017), it is also not altogether clear what visual elements LADs should possess, that is, what type of information is effective at triggering positive behavioral adjustments in learners, or what aspects are detrimental by potentially inducing anxiety among learners. Overall, evidence about whether learning analytics and LADs improve learning in practice is scarce (Ferguson & Clow, 2017; Guzmán-Valenzuela et al., 2021; Knight et al., 2020; Rets et al., 2021; Wilson et al., 2017) and has room for further investigation.

This paper expands on the challenges of operationalizing LADs and identifies gaps in current LADs. It does this by integrating analyses of recently published LADs by following a systematic line of inquiry. The following subsections offer definitions of various analytics layers that we use in our analysis of the current state of practice in LADs, as well as an outline of our research agenda for establishing an institutional-level dashboard that offers more value to learners.

Analytics layers

Analytics can provide different levels of informational insights to enable users in making informed decisions. At the most basic level, descriptive analytics highlight snapshots of variables of interest. These convey information about trends and the current

Table 1 Reviewed papers overview

Analysis		Studies
Descriptive analytics content	Conducted	Aljohani et al., 2019; Baneres et al., 2019; Bodily et al., 2018; Chatti et al., 2020; Chen et al., 2019; Fleur et al., 2020; Gras et al., 2020; Han et al., 2021; He et al., 2019; Karaoglan Yilmaz & Yilmaz, 2020; Kia et al., 2020; Kokoç & Altun, 2021; Majumdar et al., 2019; Naranjo et al., 2019; Owatari et al., 2020; Ulfa et al., 2019; Valle et al., 2021
Predictive analytics content and reported accuracy	Not conducted	Aljohani et al., 2019; Bodily et al., 2018; Chatti et al., 2020; Chen et al., 2019; Gras et al., 2020; Han et al., 2021; He et al., 2019; Karaoglan Yilmaz & Yilmaz, 2020; Kia et al., 2020; Majumdar et al., 2019; Naranjo et al., 2019; Owatari et al., 2020; Ulfa et al., 2019
	Conducted	Baneres et al., 2019; Fleur et al., 2020; Kokoç & Altun, 2021; Valle et al., 2021
	Accuracy not reported	Fleur et al., 2020; Valle et al., 2021
	80–89% accuracy achieved	Kokoç & Altun, 2021
Prescriptive analytics content	90–95% accuracy achieved	Baneres et al., 2019
	Conducted	None
	Conducted non-data driven	Baneres et al., 2019; Bodily et al., 2018; Gras et al., 2020; Han et al., 2021; Karaoglan Yilmaz & Yilmaz, 2020; Majumdar et al., 2019
	Not conducted	Aljohani et al., 2019; Chatti et al., 2020; Chen et al., 2019; Fleur et al., 2020; He et al., 2019; Kia et al., 2020; Kokoç & Altun, 2021; Naranjo et al., 2019; Owatari et al., 2020; Ulfa et al., 2019; Valle et al., 2021
Model interpretability and explainability	Conducted	None
Dashboard evaluation and effectiveness	Evaluation conducted within a pilot study context	Bodily et al., 2018; Chatti et al., 2020; Gras et al., 2020; Han et al., 2021; He et al., 2019; Kia et al., 2020; Kokoç & Altun, 2021; Naranjo et al., 2019; Owatari et al., 2020; Ulfa et al., 2019
	No evaluation conducted within a prototype study context	Chen et al., 2019; Majumdar et al., 2019
	Positive effects on student outcomes reported	Aljohani et al., 2019; Fleur et al., 2020; Han et al., 2021; Kokoç & Altun, 2021
Dashboard Color content	1-3 colors	Fleur et al., 2020
	4-6 colors	Bodily et al., 2018; Kia et al., 2020; Naranjo et al., 2019; Ulfa et al., 2019; Valle et al., 2021
	> 6 colors	Aljohani et al., 2019; Baneres et al., 2019; Chatti et al., 2020; Chen et al., 2019; Gras et al., 2020; Han et al., 2021; He et al., 2019; Karaoglan Yilmaz & Yilmaz, 2020; Kokoç & Altun, 2021; Majumdar et al., 2019; Owatari et al., 2020

status relative to other identified measures. Descriptive analytics are the simplest form of insights to extract from data, and while useful, have a limited utility.

Predictive analytics on the other hand emphasize some form of forecasting and embody the ability to estimate future outcomes based on current and past data patterns. Predictive analytics are mostly driven by machine learning algorithms which learn from historic datasets in order to produce classifiers that can make inferences about possible future outcomes from current data inputs. Data products based on predictive analytics represent a considerable increase in complexity over mere descriptive analytics and offer more value, but they also possess shortcomings. One shortcoming is that they usually produce black-box models which lack transparency into their internal workings (Adadi & Berrada, 2018). This means that it is often not possible for users to understand how these models make predictions, and what aspects of the learners' behaviors are driving the predictions towards prognosticated outcomes. This lack of model *interpretability* and *explainability* of outputs associated with most predictive models lowers their utility, and over time erodes the trust of users (Baneres et al., 2021). Therefore, a trend is emerging at regulatory¹ levels requiring predictive models to expose their reasoning behind the predictions in comprehensible ways.

The most complex and arguably the most insight-rich form of analytics is prescriptive analytics. Prescriptive analytics can leverage predictive analytics in such a way that the underlying models are also able to infer possible causal relationships and consequently generate recommendations and suggestions to users about which specific behavioral changes are most likely to result in positive outcomes. These prescriptive outputs are tailored to each learner, but their suggestions are *data-driven* and thus based on similar students who achieved positive outcomes in the past. By issuing advice on behavioral adjustments and learning strategies that learners can undertake to maximize their learning outcomes, the decision-making process for the learners can be simplified and the guesswork removed.

Currently, descriptive LADs are most commonly in use with an increasing number integrating predictive components. However, to the best of our knowledge, examples of dashboards incorporating data-driven prescriptive aspects of analytics do not exist.

Aims

This paper has three parts and contributions. We first provide an overview of the existing challenges in developing institutional LADs. We highlight three challenges (namely, representation and actions, ethics, and agility) faced in deploying LAD initiatives involving student-facing dashboards. Secondly, we provide an extensive survey of the most recently published state-of-the-art LADs in literature. Our search identified 17 LADs and we assess their common characteristics as well as strengths and weaknesses. Thirdly, we propose our LAD which addresses many of the shortcomings, and to the best of our knowledge, is the first LAD that brings descriptive, predictive and data-driven prescriptive analytics into one display. We conclude with offering inferences on what we see

¹ The General Data Protection Regulation (GDPR) is an example of a regulation requiring a "right to explanation" in respect to predictive models.

as being future directions and emerging frontiers in LA dashboarding in the short to medium term.

Our research questions are as follows:

RQ1. What are unique challenges in developing student-facing LADs?

RQ2. How ubiquitous are LADs?

RQ3. What is the current evidence for the effectiveness of LADs?

RQ4. What are the strengths and weaknesses of the current approaches to LA dashboarding?

RQ5. What are future directions of LA dashboarding and how can some of the existing weakness be addressed?

Analytics in education

There are broadly three streams of research within educational analytics. Learning Analytics (LA) focuses on learners. Its primary concern is optimizing teaching and learning processes. Educational Data Mining (EDM) on the other hand seeks to develop methods for exploring educational data in order to better understand learners, and to extract insights about them and the educational systems. Academic Analytics (AA) draws on the insights gained from educational data for supporting strategic decision-making, resolving academic issues such as retention and improving marketing strategies.

These three streams intersect at various points and share much of the underlying data, and though they could all be grouped under the same umbrella as Educational Data Science, they differ in the stakeholders which they target. EDM tends to target both teachers and learners, while LA primarily addresses the needs of learners. However, institutional administrators, managers and educational policymakers are the key stakeholders of AA applications. The three streams also affect different levels of the educational systems. LA is linked to course-level granularity and to department-level concerns within institutions, while EDM spans departmental through to faculty and institutional-level concerns (Nguyen et al., 2020). Meanwhile, AA affects universities at the institutional level, that has implications for policy making, thus it spans regional, national and possibly international levels.

Challenges for building LADs

While there are some differences between LA, AA and EDM, they all share some common challenges. Numerous studies have reported implementation details of LA products; however, a recent study by Leitner et al. (2020) pointed out that they rarely provide comprehensive descriptions of challenges faced in productionizing these systems. This study shortlisted seven general challenges for deploying LA initiatives:

1. Purpose and Gain: managing expectations of different stakeholders.
2. Representation and Actions: facilitation of actionable insights by LA products.
3. Data: communication to students regarding what is being done with their data, and formulating suitable policies to manage data processes.

4. IT Infrastructure: balancing the pros and cons of opting to use internal or external service providers for implementing and running the LA products.
5. Development and Operation: planning and implementation of the process of developing and operating an LA initiative.
6. Privacy: ensuring both security of learners' data and compliance with increasingly stringent legal requirements worldwide.
7. Ethics: ensuring that LA products do not bring harm and provide learners with the ability to opt-out.

The above challenges are generic and broadly applicable to all LA projects. We draw on recent literature to expand on two particular challenges above (2 and 7), and we tailor them to the difficulties which specifically relate to LAD projects. In addition, with supporting literature we posit an additional challenge, namely Agility, to the original seven identified by Leitner et al. (2020).

Representation and actions

Dashboard visualization is more of a science than art. The dashboard designer must possess a degree of understanding of how the human visual cortex perceives various visual cues in order to optimally match different data types to suitable visual representations. Some data are quantitative and others are ordinal or categorical in their attributes. The values of each data type are best represented by different cues which could comprise contrasting colors, differing spatial positions or variations in symbols denoting length, size, shape and orientation amongst others. The designer also needs to possess both domain expertise in learning theories and paradigms, as well as technical capabilities in developing dashboards (Klerkx et al., 2017).

Choosing the correct visualization technique can present difficulties largely due to the increasing amounts of available data and the candidate variables/indicators that can be incorporated (Leitner et al., 2019). Ensuring that dashboards are informative without overwhelming the user is a challenging balancing act. From an aesthetic perspective, Tufte (2001) cautions against use of 'non-data-ink' and 'chartjunk' in graphs, that is, he maintains that excessive use of colors, patterns or gridlines can confuse and clog the recipient's comprehension. Bera (2016) specifically mentions the overuse and misuse of color in business dashboards and the role this has on the users' decision-making abilities. Bera's research finds that contrasting colors vie for user's attention, and unless necessary, they distract and affect the decision-making processes. By using eye tracking technology, the study demonstrated that the cognitive overload associated with the misuse of color in dashboards leads to longer fixation periods on irrelevant aspects of dashboards and prolongs the ability of users to comprehend the information.

Use of predictive modelling is becoming more prominent within LA (Bergner, 2017), and these techniques are emerging more frequently within dashboards. A recent study (Baneres et al., 2021) into developing LA technologies acting as early warning systems for identifying at-risk learners highlighted the need to move beyond 'old-fashioned' dashboards that simply rely on descriptive analytics and to instead, orient efforts towards incorporating predictive analytics amongst other advanced features. However, building highly accurate and reliable predictive models is not trivial. Firstly, it requires

considerable technical expertise which is not always easy to acquire. Secondly, predicting outcomes based on human behavior reflects a non-deterministic problem. Further, for scalability reasons, we ideally require generic predictive models which can predict student outcomes across widely disparate courses. However, since courses have different attributes, styles of delivery and assessment types, it is a considerable challenge to create single generic predictive models that can work optimally across diverse courses. On the other hand, developing tailored predictive models for each different course creates technical resource overheads. Tailored models are also likely to perform badly in many instances due to scarcity of data leading to overfitting, since individual courses may have small class numbers or have limited historical data. A recent systematic literature review on the current state of prediction of student performance within LA, Namoun and Alshanjiti (2020) found that the state of predictive modeling of student outcomes is not fully exploited and warrants further work. The study found that not only the accuracy of the existing models has room for improvement, but more robust testing for its validity, portability (or generic models) and overall generalizability needs to be conducted. In a recent study, Umer et al. (2021) concluded that many datasets used to build predictive models in this domain were small, often having less than 10% of the overall data points for certain class labels, leading to unreliable predictive accuracies especially when course-tailored predictive models are being created. The study also calls for enhancing the scope of engagement data to cover learner interaction data from forum messages, identifying pedagogically meaningful features and developing dashboard visualizations that have some underlying pedagogical intent.

Developing accurate classifiers is further complicated by the negative effects of concept drift (Lu et al., 2018). Concept drift describes the degradation in accuracies of predictive models over time since data used to build models may become disconnected with current real-life data. This can occur when learners' study patterns gradually or abruptly change (as in the case of pandemic responses), and current digital footprints no longer correlate with previous patterns in the historic record. For example, the gradual shift towards the use of virtual learning environments (VLE) over the last 10–15 years represents a concept drift. Learners' study patterns prior to this period in the historic record bear little resemblance to the patterns of learners of today, and thus, data from historic period will likely degrade predictive accuracies of current students. Concept drift can also happen suddenly, as indeed the sudden migration to full online learning during the recent pandemic crisis brought into play additional technologies and different digital patterns and footprints that students leave behind. This disconnect between the independent and dependent variables from historic data needed to train the predictive models, with the independent variables being used as input to predict the outcomes of current students, is constantly evolving. This phenomenon represents a technical and a capability challenge for universities, as concept drift needs to be detected and accounted for, while the mechanisms for achieving this effectively are still being researched (Lu et al., 2018).

The above challenges are considerable. However, even if they can all be addressed, it is now no longer sufficient to deploy predictive models and solely display their outputs without providing the learners with explainability of how a model arrived at a given prediction. It is also becoming more apparent that learners will engage with a LAD only if

they understand how displayed values are generated (Rets et al., 2021). Liu and Koedinger (2017) argue for the importance of interpretability which leads onto actionability. Models need to possess explanatory characteristics so that learners understand why a model produced given predictions, what the underlying driving factors are, and importantly, what insights can be derived from these explanations in order to trigger actionable behavioral adjustments. Not only should interpretability of models and explainability of their individual predictions be provided to the learners, but also counterfactuals, which explicitly demonstrate alternative outcomes for the learner if a behavioral change were to take place in specific areas. Recent studies (Rets et al., 2021; Valle et al., 2021) in LADs have highlighted the necessity of integrating insights which are prescriptive and take on forms of recommendations to guide students in their learning. Producing such rich and sophisticated outputs is a challenge, because extracting simplified representations of predictive black-box models and their reasoning is complex. There are limited available tools with sufficient maturity that support this functionality, which again requires a high level of expertise to implement and leverage.

Ethics

The challenges surrounding ethical use of data within LA products are generally well understood and accepted. They center around questions of what personal data should be collected and processed by these systems, what insights should be extracted and with whom they should be shared. Additional concerns exist around possible consequences on learners when conveying personalized information; therefore, institutions need to be aware of intrusive advising or inappropriate labelling that may lead to learner resentment or demotivation (Campbell et al., 2007). As such, avoidance of harm to learners, alongside compliance with legal requirements are paramount.

Given the importance of practical ethical underpinnings when using LA systems, it is acknowledged that robust and clear policies need to be formulated on what empirical data is permitted to be used for analytical purposes and to what end (Kitto & Knight, 2019). The study supports that awareness of these policies must be communicated to the learners together with the purported educational benefits that such systems claim to bring, together with the potential risks. A key concern however is the uncertainty regarding the benefits distribution, which may not be the same for everyone (Rubel & Jones, 2016); hence, institutions are encouraged to create a sense of transparency about LA systems by including statements on their data practices and limitations.

Beyond the well accepted dilemmas of LA systems listed above, predictive models used in LADs bring with them some other acute challenges. Predictive models naturally embody within them the process of generalization. As the machine learning algorithms learn and induce predictive models, they move from individual and specific examples to more general descriptors of the data. With this natural induction process, errors are invariably introduced. The ethical concern and challenge come into play when we consider both incorrect and correct classifications and the effects that they might have on learners. If a student is mis-classified as being “at-risk” this might have the effect of discouraging them and eventuate in the “fulfillment of the prophecy” despite the fact they were originally on-track to successful completions. Or, in using oversimplified classification labels, we can diminish the predictive value and in turn reduce the trustworthiness

of the analytical approach. This challenge will always remain since learners are not deterministic and predictive models in non-deterministic domains are inherently imperfect. Likewise, Bowker and Star (2000) note that even with correct predictions, for some this may be an incentive if they are already motivated and capable of positively adjusting their course in order to alter their predicted outcome, while for others, the prediction may only serve to further deflate.

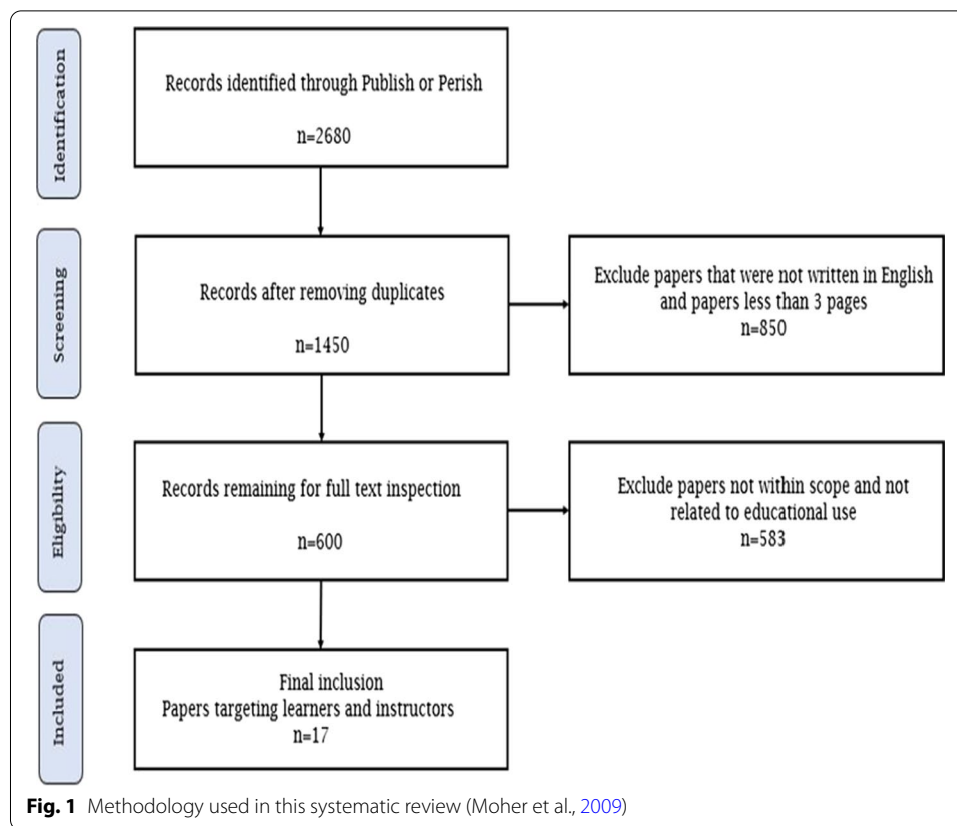
Agility

Agility is the ability to rapidly adapt to changing requirements, be flexible and able to seize new opportunities. Universities are more resistant to change than industrial entities (Menon & Suresh, 2020); they are typically considered to be fractured and decentralized (Bunton, 2017), while possessing complex and non-standard business processes (Mukerjee, 2014a). However, financial constraints coupled with pressure from competition as a consequence of the unfolding digital revolution, have put universities on high alert to engage with new technologies (Mukerjee, 2014a). It is recognized that organizational agility is a crucial capability for universities at these times (Mukerjee, 2014b). Both the use of data insights and analytics as well as the development of these projects, places immediate demands of agility on behalf of the organization operationalizing them. Agility is therefore a key challenge for universities attempting to productionize LADs.

The requirement for agility comes at different levels in respect to LADs. Translating LADs into products that genuinely improve learning outcomes requires constant monitoring and analysis of their usage patterns, user feedback and ultimately the gathering of evidence into their efficacy. The consequences of this are an increase in resource costs for maintenance and continuous refinement of the LADs. Continuing support from the institutions and willingness to provide ongoing long-term refinements need to be secured ahead of time. Sun et al. (2019) point out that improvements of these types of systems needs to go beyond pilot and deployment stages, and that underlying assumptions used to develop these systems need to be re-assessed as adjustments are made to enhance the design or functionality. For best results, the design of dashboards should be iterative with continuous feedback from learners in order to ensure that an operationalized product is actually useful. This is time and resource intensive and requires agility.

From a data-oriented point of view, agility and the ability to integrate new data streams into LADs are paramount. Universities are rapidly incorporating modern technologies for course delivery and improving the learning experience. The technologies sometimes augment what is already in place, while other times, they completely replace legacy processes and systems with new ones. This process has been accelerating recently and will continue to do so. The consequence is that new and more diverse digital footprints will continue to be generated by learners especially with the increased demand in online education in the aftermath of COVID-19. Therefore, adaptability and rapid responses in integrating new data sources must be set forth to identify new features that can improve the predictive power of deployed models.

Finally, profound insights are compelling. They demand action if negligence is to be avoided. Deep insights can be game-changers and often call for swift action even when this is inconvenient. For example, if predictive models powering the LADs identify certain qualifications within an institution's portfolio as being key predictive drivers towards



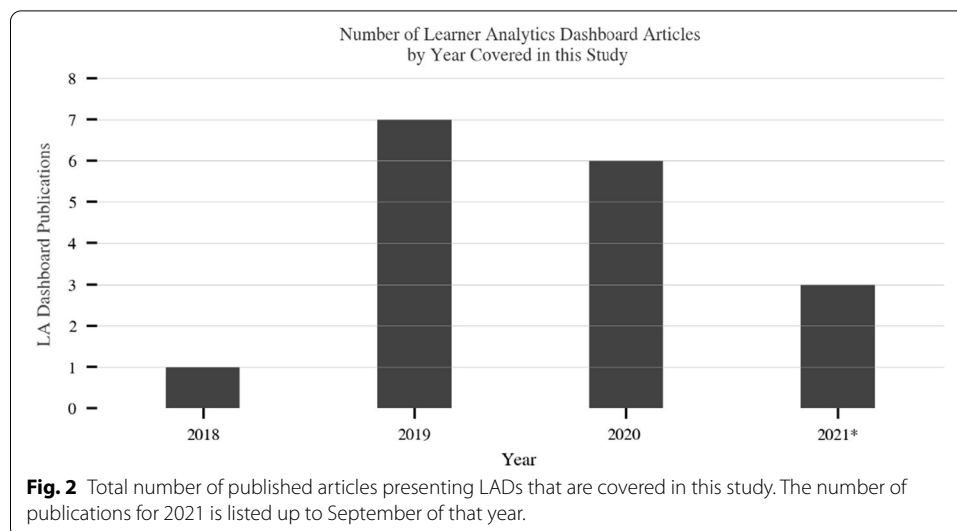
poor completion rates, then this would need to trigger action and possibly advice on changes that may neither be convenient for an institution, nor even align with their overarching strategic goals. With deployment of LADs, therefore, comes the responsibility of asking the tough questions in adapting to the suggested changes that can have a better institutional impact.

Methods

The focus of this study was to review the most recent developments in LADs. To that end, the search focused on studies published from 2018 until the time the search was completed (September 2021). The search followed the PRISMA framework (Moher et al., 2009) which requires a principled approach to defining the inclusion and exclusion criteria as well as search parameters.

We first conducted a keyword search targeting Google Scholar using the “Publish or Perish” tool in order to retrieve the initial academic articles. The following search terms were used: “learning analytics dashboard” or “visualization tool” or “early warning system” or “student dashboard”. The search yielded the following total number of results per year: 2018 $n=340$, 2019 $n=977$, 2020 $n=960$, and 2021 $n=403$. A total of 2680 papers were obtained. This was reduced to 1450 papers following the elimination of duplicates.

Next, papers that were not written in English and those containing less than 3 pages were filtered, resulting in 600 papers. The abstracts of these papers were screened, and finally only the papers that focused on dashboards targeting learners and instructors



were retained. This yielded a total of 17 papers that successfully passed all the inclusion criteria and only these were included in the final analysis. Figure 1 outlines the overall methodology used for data collection, while Fig. 2 depicts the histogram of the 17 LAD papers by year.

Against the backdrop of recently published literature, the first part of our study has already identified challenges facing LA and specifically difficulties associated with the development and deployment of dashboards in educational contexts. The second part of our study analyzes the data on existing LADs. Our analysis approach is based on five key assertions that are grounded in dashboard literature. The assertions act as a prism through which we reviewed the LADs and directed our investigation towards the design of our dashboard subsequently. These are as follows:

1. Given that LADs using only descriptive analytics is not enough, it is meaningful to identify dashboards which have started to incorporate predictive and data-driven prescriptive analytics.
2. Since accuracy of predictive models is an identified challenge, it is informative to determine what accuracies are being reported for recent LADs using predictive modeling, and if they communicate the confidence of their predictive outputs to the learners.
3. Assuming that there is value in providing learners with some level of interpretability of the underlying predictive models and explanations of how the models have arrived at predictions for individual students, it is instructive for future research directions to ascertain how prevalent is the presentation of these features on the LADs.
4. Since the evidence of the effectiveness of LADs to affect positive outcomes for learners is not complete, it is instructive to know what evaluation efforts have been made in recent studies on the utility of LADs.
5. LADs using higher number of colors are more likely to misuse and overuse colors and contribute towards confusion for the learners.

Analysis

Our analysis is divided into two parts. The first part reviews each of the 17 LADs and highlights noteworthy aspects of each one. The second part analyzes the dashboards at an aggregate level and offers analyses of the state of LA dashboarding in a summative format.

Review of dashboards from literature

Our framework for reviewing all the LAD studies uses a scheme whereby we consider each LAD from the perspective of how they have implemented descriptive, predictive and analytics functionalities, as well as the reported evidence outlining the effectiveness of LADs on learner outcomes.

LADs with descriptive analytics capabilities

All studies incorporated some aspects of descriptive analytics. Frequently, the descriptive analytics were in the form of graphical displays depicting comparisons of a student in respect to class averages or patterns in relation to students (Aljohani et al., 2019) across metrics like assessment scores, participation levels and interaction with online activities (Chen et al., 2019; Fleur et al., 2020; Gras et al., 2020; Han et al., 2021; Karaoglan Yilmaz & Yilmaz, 2020; Kokoç & Altun, 2021; Ulfa et al., 2019; Valle et al., 2021). Some studies focused on status updates of progression through online course materials such as video, reports on time spent on eBooks and summaries of course notes (Bodily et al., 2018; He et al., 2019; Majumdar et al., 2019; Owatari et al., 2020). Other studies also assisted learners in planning and provided alerts of upcoming assessment submission deadlines (Baneres et al., 2019; Kia et al., 2020; Naranjo et al., 2019). Certain LADs (Chatti et al., 2020) went beyond static dashboards and enabled direct customizations of them by allowing learners to dynamically generate indicators of their choice.

LADs with predictive analytics capabilities

Several studies went further than mere descriptive analytics and incorporated predictive analytics elements into their dashboards. A descriptive and a predictive dashboard was developed by Valle et al. (2021). The descriptive dashboard aimed at displaying the students' performance relative to the class average while the predictive dashboard displayed the probability of learners attaining specific grades. The authors reported that the predictive dashboard helped only the highly motivated students to sustain their motivation levels, while both dashboards failed to demonstrate their effectiveness in affecting final outcomes. Similarly, Fleur et al. (2020) developed a LAD with class-comparative descriptive components as well as the student's predicted final grade. The students in the treatment group accessed the dashboard and their performance was analyzed in formative and summative assessments. The study reported that students in the treatment group performed better in the formative assessment only. Baneres et al. (2019) focused on devising an early warning system for learners and instructors that identifies at-risk students. Their Graduate At-Risk (GAR) model used grades to predict course outcomes. Additionally, an intervention

mechanism was incorporated that automated sending personalized messages to at-risk students. While GAR noted an improvement in performance of the at-risk students, it could not be determined which factors were responsible. In a similar vein, a prescriptive learning dashboard (PLD) using personalized recommendation texts was developed by Kokoc and Altun (2021) which also focused on generating of student risk status and displaying it on the dashboard. The authors concluded that those students who used PLD performed significantly better in their courses. None of the examined studies took steps to communicate to students through the dashboard how reliable the underlying predictive models were, nor were technologies used which could elucidate to learners how the models operated, or how the predictions were generated based on specific student's data.

LADs with prescriptive analytics capabilities

Certain studies already mentioned above like Kokoc and Altun (2021) and Baneres et al. (2019), considered the dispatch of personalized messages as prescriptive components. Indeed, a number of other studies also leveraged different forms of messaging, recommendation techniques and communication features through the LADs in order to claim prescriptive capabilities. Bodily et al. (2018) developed LADs that recommended content and skill-building activities such as practice exercises. Their study noted that skill-related recommendation components were found by students to be more useful compared to the content recommender features. Along similar lines, Karaoglan Yilmaz and Yilmaz (2020) took the approach of delivering weekly reports over the course duration along with personalized recommendations to each student. The study claimed that providing analytics reports positively increased student motivation. Gras et al. (2020) expanded the capabilities of their LAD by providing students with an action button which accessed direct help from the instructor and can therefore be categorized as having prescriptive aspects. Direct contact between students and instructors was also enabled by the face-to-face collaborative argumentation (FCA) dashboard developed by Han et al. (2021). This tool monitored students' learning progress and facilitated prescriptive interventions by instructors with students requiring additional assistance. Meanwhile, LAView a dashboard was developed by Majumdar et al. (2019) which computed an engagement score as an aggregate value across several student interaction measures. Based on the engagement score, the instructor initiated prescriptive measures in the form of personalized emails to corresponding students. LADs are clearly emerging with some forms of prescriptive components, though many of these can also be defined as human interventions. Others which have more of an automated algorithmic approach to dispensing recommendations of content and activities are based on simplistic hard-code heuristics and thresholds. More sophisticated prescriptive components within LADs leveraging algorithmic and *data-driven* analytics have yet to emerge.

Reported effectiveness of LADs

Value assessments of the various LAD projects have taken two distinct approaches amongst the examined studies. Some studies (e.g., Bodily et al., 2018; Chatti et al., 2020; Gras et al., 2020; Han et al., 2021; He et al., 2019; Kia et al., 2020; Kokoç & Altun, 2021; Naranjo et al., 2019; Owatari et al., 2020; Ulfa et al., 2019) have largely conducted

qualitative evaluations of the various LAD deployments within pilot contexts. These included surveys and interviews of usability aspects and covered subjective responses on the degree that LADs facilitated learning. However, other studies provided quantitative findings supported by statistical analyses that demonstrated that LAD usage had positive effects on student outcomes (Aljohani et al., 2019; Fleur et al., 2020; Han et al., 2021; Karaoglan Yilmaz & Yilmaz, 2020; Kokoç & Altun, 2021).

Dashboard data analysis

Table 1 summarizes all the revised dashboards through the analysis approach listed in “Methods” Section. We find that 59% LADs include some form of descriptive analytics information to the learners. The remaining LADs focus on assisting students with planning, helping them monitor progress through online learning materials and provide learners with a medium through which instructors can more effectively interact with the learners.

The majority of the LADs either did not use any form of predictive analytics or did not report on this capability if implemented. This large group consisted of 76% of the most recently developed dashboards for learner-facing educational contexts. Out of the remaining 24% which did use predictive analytics, one of the dashboards generated predictive models with an accuracy range between 80 and 89%, while another achieved higher accuracies reaching up to 95%; however, while these accuracies were reported in literature, the model accuracies were not presented to the students on the dashboards. The remaining dashboards that used predictive analytics did not report on their predictive accuracies.

The data also indicates that model transparency approaches and technologies have not entered usage amongst the dashboard developers. Of all the dashboards which used predictive modeling, we find that no attempt was made to offer model interpretability to the learners in terms of what were the key features. Additionally, we find that the none of the reviewed dashboards tried to explain to the learners how the predictive models actually arrived at the predictions that were presented to them.

From our review, we found that none of the recent LADs utilize data-driven prescriptive analytics. Our data indicates that 47% used some form of prescriptive features associated with the dashboards which took the form of encouraging messages, supportive emails or instructive suggestions being issued to learners by the teachers. However, none employed automated instructions or recommendations generated by prescriptive modelling algorithms.

Widely contrasting approaches were adopted by researchers in respect to evaluating the usability of the dashboards, as well as their overall ability to affect positive learning outcomes. We found that 59% of the dashboards which were deployed in some form of productionized environment, evaluated the usability of the dashboards through qualitative approaches that involved surveys and interviews with learners. A further 12% of the studies created dashboard prototypes and conducted a qualitative investigation into their usability, while a same proportion developed prototypes, but did not evaluate them. A quarter of the studies performed a qualitative investigation into the effectiveness of the dashboards ability to impact student outcomes. These studies concluded that their LADs exhibited a positive impact on student outcomes. A common feature across

Table 2 Dashboard technologies and size of study cohorts

Study	Technology	Programming/ expertise	Cohort size
Bodily et al., 2018	N/A		180
Chen et al., 2019	N/A		–
Aljohani et al., 2019	ASP MVC4, HTML5, JQuery and High-charts JavaScript	High	86
Ulfa et al., 2019	N/A		67
Majumdar et al., 2019	N/A		–
He et al., 2019	HTML5, JavaScript and Echarts	High	327
Naranjo et al., 2019	Vue.js, HTML, CSS	High	64
Baneres et al., 2019	Web application	High	247
Gras et al., 2020	N/A		127
Karaoglan Yilmaz & Yilmaz, 2020	LMS messaging tool	Low	81
Fleur et al., 2020	Django	High	79
Chatti et al., 2020	Google charts and C3.js	High	414
Kia et al., 2020	JavaScript, D3.js	High	449
Owatari et al., 2020	Web application	High	108
Han et al., 2021	Web application	High	88
Kokoç & Altun, 2021	Google visualization API and AJAX API	High	126
Valle et al., 2021	R and Shiny	High	179

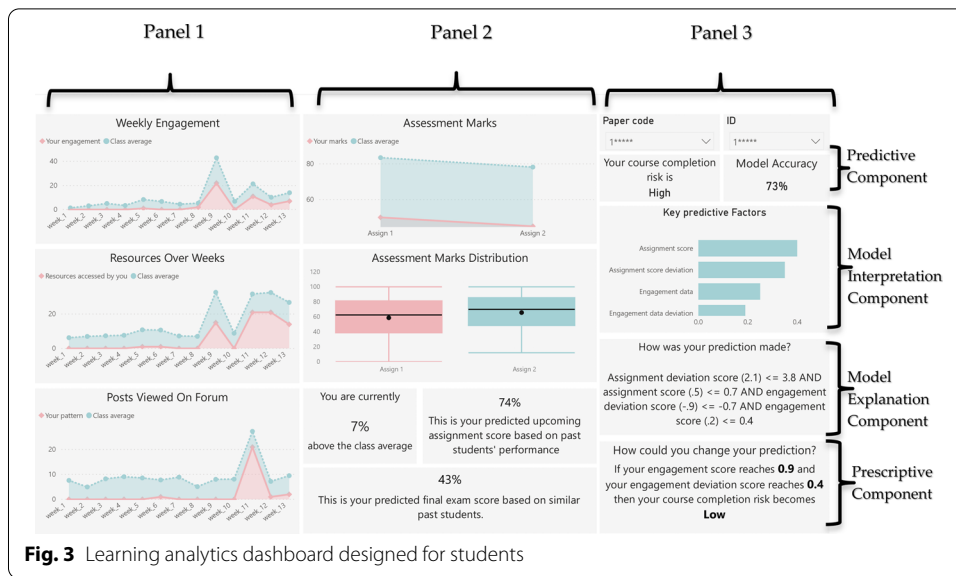
half of the LADs which demonstrated a positive student outcome was they possessed predictive analytics capabilities. They all depicted information for each student in relation to where they were situated in respect to their peers across various metrics, and one of the dashboards implemented prescriptive features.

Our analysis indicates that 59% of the LADs have used six or more different colors on their display, potentially contributing towards information overload and miscommunication of insights. 29% used between four and six colors, while the remaining 12% employed up to 3 colors only.

Technology is an important aspect through which LADs ought to be considered. The chosen technology determines the range of capabilities of LADs and the agility of the projects. Where reported, Table 2 indicates that the chosen tools for implementing LADs have so far mostly been web application frameworks which carry with them a requirement of a high level of technical expertise. Usage of off-the-shelf commercial dashboarding products which do not require a high level of technical and programming expertise appear not to be a chosen medium yet. Both web application frameworks and off-the-shelf dashboarding products generally possess very limited advanced analytics capabilities, which would then require additional technologies to overcome this limitation. The exceptions being, Shiny (R) and Django (Python), where both technologies have access to a large ecosystem of analytics capabilities.

Proposed learning analytics dashboard

The previous section reviewed dashboards and highlighted their key themes, analytic capabilities and reported efficacies. Drawing on contributions from these studies as well as the strengths and weaknesses of various dashboarding approaches, we propose our learner-facing dashboard design (shown in Fig. 3). The proposed dashboard



attempts to integrate all levels of analytics capabilities missing in reviewed dashboards. The proposed dashboard has descriptive, predictive as well as prescriptive components built into it. To the best of our knowledge, this dashboard is the first of its kind to embed data-driven prescriptive capabilities involving counterfactuals into its display. In addition, our dashboard possesses a high degree of transparency and communicates to the learners how reliable the predictive models are; what the key factors are that drive the predictions, as well as the conversion of a black-box predictive model into a glass-box, human interpretable model for the learners so that they can understand how their prediction is being derived.

Analytics layers

The proposed dashboard distributes the descriptive, predictive and prescriptive components across three panels seen in Fig. 3. The first panel highlights student engagement levels. This panel contains only descriptive analytics components and compares the learner’s engagements versus that of the cohort’ average. Engagement includes weekly login counts into the virtual learning environment, number of learning resources accessed, and total forum posts created as a measure of communication exchange levels.

The second panel displays information regarding a learner’s academic performance. This panel has both descriptive and predictive analytics components. The descriptive component in the top half, displays the snapshot of a learners’ assignment grades, quizzes and tests. The learner’s data is contrasted again with that of the cohort. The student can rapidly see their deviation from the class mean and can also inspect in greater detail how far their score deviates from their peers by viewing the overall class distribution through the box-and-whisker plots. The dashboard’s predictive component begins in the lower half of the second panel. This component provides a student with estimates of what scores they are likely to achieve in the upcoming assignment and their final exam based on the learners who have exhibited similar learning attributes in the past.

However, the key predictive analytics component and the novel prescriptive analytics features are found in the third panel. In this panel, an overall prediction is made regarding the learner's estimated risk profile for meeting the course's learning outcomes. Given the importance of this model, we emphasize key aspects of its nature that were missing in previous studies. The dashboard communicates the accuracy of the underlying model to the learner, and provides interpretability to the user in terms of what factors are deemed important to the model at a high-level when it makes a prediction. In addition, the dashboard contains an explainability component which communicates to the learner how the model has arrived at a given prediction for their individual case with the student's specific input values. The model reasoning provides the learners with a suggestion of what they can alter in their learning behavior in order to alter their outcomes.

The above model transparency capability is further built upon and expanded by the dashboard's prescriptive analytics features which incorporate counterfactuals. The counterfactuals indicate to the learner what specific factors together with minimal changes to their values, would produce different, and more positive predictive outcomes. The counterfactuals make some plausible assumptions about the existence of causal links in the underlying data, and based on this, generate automated advice to learners about how to maximize their learning outcomes.

From an aesthetic point of view, the dashboard attempts to minimize the use of color and renders the display in three hues, thus minimizing the risk of information overload. Additionally, the dashboard uses a neutral pastel palate to further reduce negative effects that colors can have, while attempting to maximize the data-to-ink ratio.

From a functional point of view, the proposed dashboard provides comprehensive analytics capabilities that are not found in existing LADs and demonstrates the state-of-the-art in terms of incorporating these functionalities. However, the dashboard is currently in a pilot stage at a tertiary institution with students from across 20 classes actively trialing the tool and evaluating it for usability; therefore, data on its effects on outcomes is not yet available.

Dashboard design details

The underlying data for the dashboard originated from Moodle, an open-source learning management system which provides the virtual learning platform for e-learning at the institution. The dashboard was implemented following a client server architecture. On the client-side the Power BI² tool was used to develop the web-based dashboard application. Meanwhile, on the server-side Python³ was used for both analytics and for the extract, transform, load phases.⁴

We used a mixture of Python's scikit-learn library and the CatBoost (Dorogush et al., 2018) classification algorithm for generating student outcome predictions. The underlying features used to make the predictions were engagement deviation score, engagement

² Power BI is a commercial software package owned by Microsoft. Effective at building dashboards, but at present lacks capabilities for predictive and prescriptive modelling, as well as model interpretability.

³ Python is a general-purpose programming language with a rich set of capabilities for implementing all required analytics; however, it lacks easy-to-use capabilities for building front-end dashboards.

⁴ Both Power BI and Python can effectively be substituted for R and it's R Shiny technology for constructing dashboards and developing the underlying analytics functionalities.

rolling average score from the Moodle, assignment rolling average score, assignment deviation score, and previous grade from the student management system, and demographic information such as age, English equivalent test, and highest school qualification. Models were trained on a dataset comprising 4000 students. Hold-out method was used and the final accuracy from all the test datasets was displayed to the learners as a measure of confidence in the reliability of the underlying model. The dataset was divided into weeks and were used for the prediction analysis. Moreover, the data was used in a cumulative fashion for making predictions. For example, when making predictions for week 2 the students' data from week 1 were taken into consideration as well. The reason being that prediction accuracy improves as more data becomes available in upcoming weeks. Prediction accuracy at early stages is important so that timely interventions can be made to help students.

Model interpretability was implemented used feature importance analysis which depicts the relative contribution and importance of each variable towards making predictions. This was also complemented with the use of anchors⁵ (Ribeiro et al., 2018). Anchors have recently devised as an approach for making black-box models interpretable. Anchors create proxy models which mimic the behavior of the underlying black-box model but present themselves to the user as a glass-box model. Proxy models are approximations of the real model and they present themselves as succinct human-readable decision rules.

In order to realize prescriptive capabilities, we used data-driven counterfactuals (Wachter et al., 2017) to suggest to students how an adjustment in certain behavioral learning patterns would result in a more positive prediction. For example, the counterfactual may suggest to a learner that an increase in their next assignment mark by a specific amount would change their classification from high-risk to low-risk. Such data-driven counterfactuals are based on correlation and do not guarantee causal links; however, in many cases when features are judiciously selected some degree of potential causality can safely be assumed. We use the Python counterfactual library⁶ (Mothilal et al., 2020) to generate the prescriptive analytics on the dashboards. The advantage with this tool is that the outputs are once again in a rule-based format and easy to comprehend. Additionally, the prescriptive suggestions represent a minimum shift in the values of key features that would need to take place in order to achieve a different outcome to what is currently predicted.

Discussion

Our study reveals that learner-facing LADs are steadily gaining popularity (Fig. 2), while it is reasonable to assume that numerous others may have been deployed but remain unpublished. While the value of LADs are recognized by education providers, we find from literature that many of the published dashboards are only in their prototype phases, and only few in the pilot implementation stages. This is also in agreement with findings from other studies (Chen et al., 2019; Karaoglan Yilmaz & Yilmaz, 2020). A speculative link could be argued between the low deployment rates of LADs covered in this study

⁵ Our implementation of anchors used the anchor-exp library <https://pypi.org/project/anchor-exp/>

⁶ Dice-ml <https://pypi.org/project/dice-ml/>

and the underlying technology choices taken as seen in Table 2. The technologies used in the studies are heavy-duty in respect to design, development and maintenance of LADs, requiring significant resource investments and agility. As discussed, higher education organizations are short on both of the latter requirements at present which are therefore possibly contributing factors. Given that the reviewed LADs mostly used only descriptive analytics, it could be argued that off-the-shelf commercial dashboarding software would have delivered the same functionalities for a vastly reduced effort, with higher prospects of productionization.

Given the significant resources required to operationalize LADs, our study has revealed that there is paucity of evidence on their effectiveness to affect learner outcomes. This is again supported by Fig. 2 which suggests that learner-facing LADs are a relatively new and emerging technology in the LA space and so comprehensive and conclusive meta-study research into their effects has simply not yet taken place. We also see in Table 2 that most of the past LAD papers involved relatively small study cohorts to support conclusive findings, with the median being 126 subjects. Larger studies involving several hundred subjects are emerging, and more will be needed in future in order to answer this question concretely. Encouragingly, our research did find that about a quarter of the studies concluded that their LADs produced positive impacts on student outcomes. However, the number of studies were too small to determine which types of visualizations or dashboard features directly contributed towards positive impacts on learner outcomes. Further still, it is unclear if any effects could be attributed to dashboards themselves, or to the associated human interventions.

Given the ubiquity of machine learning now, it is a little surprising that predictive modelling has not featured in a larger percentage of reviewed LADs. A possible hypothesis could be that stringent ethics requirements and risk-averse positions taken by Research Ethics Committees may be playing a role. Some of the research using predictive analytics could also be encountering obstacles due to emerging legal requirements that predictive modelling is made completely transparent, interpretable and the predictions explainable to those affected, thus lifting much higher the barrier to entry for those seeking to leverage machine learning.

Undoubtedly though, data-driven prescriptive analytics represents the next frontier of LAD development. This is the most sophisticated level of analytics with the ability to offer learners evidence-based concrete suggestions or recommendations about what adjustments in learning behaviors would most likely result in positive outcomes (Aljohani et al., 2019; Lu et al., 2018).

Future directions

Our final research question considers future directions of LA dashboarding and inquires into how some of the existing weaknesses can be addressed. We find that personalization of learning, which could be referred to as “precision learning” is the future, and there is a role for LADs in supporting this through embedding of recommendation-like features which suggest next steps to learners for maximizing outcomes. In addition, LADs can take on greater roles in early intervention responses to learners identified as being at-risk. Integrating automated interventions within the dashboards and evaluating their effectiveness will be one of the future research directions. The focus on personalized

learning and early interventions as an area of needed focus is also supported by (Gras et al., 2020; Han et al., 2021) while closing the loop and ensuring that the insights generated by LA systems, or dashboards, is actionable and not just interesting, is emphasized by (Baneres et al., 2019; Chen et al., 2019).

Maisarah et al. (2020) noted in their broad survey of LADs the importance of embedding customization capabilities within dashboards in order to make them user-friendly and thus promote long-term usage of the dashboards. Future research will focus on developing technologies that possess these capabilities and are able to seamlessly integrate with native platforms used by institutions for their existing Virtual Learning Environments. Leitner et al. (2019) also mention the utility of embedding analytics within dashboards themselves in order to directly gather information on learner usage patterns of the dashboards themselves in order to optimize them in subsequent iterations of development.

Lastly, Sedrakyan et al. (2020) go further and ambitiously suggest integrating data from activities in the learning-process which may not be directly linked with the institutional learning environments. They propose data acquisition from multi-modal sources such as biofeedback from various wearable sensors, audio/video streams and using them to augment LADs. Thus, scalability in processing capabilities of live data streams originating from wearable sensors would form yet another requirement of future work for LADs.

Study limitations

We acknowledge that the search time-window of 2018 to 2021 is constrained, and that data from 2021 is partially collected which constitutes a limitation of this study. Data on the usability of the proposed dashboard and its effects on student outcomes are being collected. A further limitation of this study is that these data cannot yet be presented, neither can this tool be made available publicly for trial purposes at this point in time due to software licensing constraints.

Conclusion

Learning Analytics dashboards (LADs) are becoming increasingly commonplace within the educational sector with the aims of improving the quality of the learning experience and thereby maximizing learner outcomes. Our study focused on identifying challenges associated with LAD projects as well as analyzing characteristics of recent advances in LADs. We comprehensively surveyed existing LADs and analyzed them through the prism of the sophistication of insights they deliver and ways in which they help learners make informed decisions about making adjustments to their learning habits. Finally, in considering the strengths and weaknesses of existing LADs, we propose a dashboard currently being deployed for trials at a tertiary institution that attempts to address some of the gaps we found in literature. Our research findings have both theoretical and practical implications.

Theoretical implications

We have added to the body of knowledge surrounding what we know to be challenges in operationalizing Learner Analytics (LA) projects. We refined these challenges to LAD projects and have identified the lack of agility in higher education institutions as

one of the key pressure points. Our work has confirmed that learner-facing LADs are on the rise within higher education institutions, but significant gaps in understanding and quantifying the effectiveness of LADs exists. In particular, uncertainty exists about which components within LADs are more effective at improving learning outcomes. We find that predictive modeling functionalities are not used in majority of cases within the reviewed LADs, and examples of interpretability of the models and the ability to explain their predictions to the learners do not yet exist in published studies. Additionally, our study reveals the absence of data-driven prescriptive analytics which, with other gaps, highlights numerous worthwhile avenues for future studies to pursue.

Practical implications

A key practical implication of this study is a demonstration of how a sophisticated LAD can be developed which integrates all forms of analytics: descriptive, predictive and prescriptive. We have demonstrated how interpretability of predictive models can be made available to the learners and critically, how the specific predictions for a given learner can be explained to them. This will establish trust with the users through transparency of moving beyond black-box predictive models, and in the process satisfy emerging regulatory requirements. Additionally, we have demonstrated how automated and data-driven prescriptive analytics can be leveraged within LADs. Our research also points the analytics practitioners towards recently developed technologies which more than ever, make these capabilities accessible to the wider audience.

Abbreviations

AA: Academic Analytics; EDM: Educational Data Mining; LA: Learning analytics; LADs: Learning analytics dashboard(s); RQ(s): Research question(s); VLE: Virtual Learning Environment.

Acknowledgements

Not applicable.

Authors' contributions

Conceptualization, TS, GR and AM; Methodology, TS and GR; Software, TS and GR; Validation, TS, GR and AM; Formal analysis, TS; Investigation, TS and GR; Data curation, TS and GR; Writing—original draft preparation, TS; Writing—review and editing, TS and AM; Visualization, TS and GR; Supervision, TS and AM; Project administration, TS. All authors read and approved the final manuscript.

Funding

Not applicable.

Availability of data and materials

Not applicable.

Declarations

Competing interests

The authors declare no competing interests.

Received: 23 October 2021 Accepted: 14 December 2021

Published online: 14 February 2022

References

- Adadi, A., & Berrada, M. (2018). Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). *IEEE Access*, 6, 52138–52160. <https://doi.org/10.1109/ACCESS.2018.2870052>
- Aljohani, N. R., Daud, A., Abbasi, R. A., Alowibdi, J. S., Basher, M., & Aslam, M. A. (2019). An integrated framework for course adapted student learning analytics dashboard. *Computers in Human Behavior*. <https://doi.org/10.1016/j.chb.2018.03.035>

- Aristovnik, A., Keržič, D., Ravšelj, D., Tomaževič, N., & Umek, L. (2020). Impacts of the COVID-19 pandemic on life of higher education students: a global perspective. *Sustainability*. <https://doi.org/10.3390/su12208438>
- Baneres, D., Guerrero-Roldán, A. E., Rodríguez-González, M. E., & Karadeniz, A. (2021). A predictive analytics infrastructure to support a trustworthy early warning system. *Applied Sciences*. <https://doi.org/10.3390/app11135781>
- Baneres, D., Rodríguez, M. E., & Serra, M. (2019). An early feedback prediction system for learners at-risk within a first-year higher education course. *IEEE Transactions on Learning Technologies*. <https://doi.org/10.1109/TLT.2019.2912167>
- Bera, P. (2016). How colors in business dashboards affect users' decision making. *Communications of the ACM*. <https://doi.org/10.1145/2818993>
- Bergner, Y. (2017). Measurement and its uses in learning analytics. *Handbook of Learning Analytics*. <https://doi.org/10.18608/hla17.003>
- Blankenberger, B., & Williams, A. M. (2020). COVID and the impact on higher education: The essential role of integrity and accountability. *Administrative Theory & Praxis*, 42(3), 404–423. <https://doi.org/10.1080/10841806.2020.1771907>
- Bodily, R., Ikahihifo, T. K., Mackley, B., & Graham, C. R. (2018). The design, development, and implementation of student-facing learning analytics dashboards. *Journal of Computing in Higher Education*, 30(3), 572–598. <https://doi.org/10.1007/s12528-018-9186-0>
- Bowker, G. C., & Star, S. L. (2000). *Sorting things out: Classification and its consequences*. MIT Press.
- Bunton, T. E. (2017). *Agility within higher education it organizations: a loosely coupled systems perspective*. <https://dc.uwm.edu/etd/1451>
- Campbell, J. P., Deblois, P. B., & Oblinger, D. G. (2007). *Academic Analytics: a new tool for a new era*.
- Chatti, M. A., Muslim, A., Guliani, M., & Guesmi, M. (2020). *The LAVA model: learning analytics meets visual analytics*. https://doi.org/10.1007/978-3-030-47392-1_5
- Chen, L., Lu, M., Goda, Y., & Yamada, M. (2019). Design of learning analytics dashboard supporting metacognition. *16th International Conference on Cognition and Exploratory Learning in Digital Age, CELDA 2019*, 175–182. https://doi.org/10.33965/celda2019_201911022
- Dorogush, A. V., Ershov, V., & Gulin, A. (2018). *CatBoost: gradient boosting with categorical features support*. <http://arxiv.org/abs/1810.11363>
- Ferguson, R., & Clow, D. (2017). Where is the evidence? A call to action for learning analytics. *Proceedings of the Seventh International Learning Analytics & Knowledge Conference*, 56–65.
- Fleur, D. S., van den Bos, W., & Bredeweg, B. (2020). *Learning analytics dashboard for motivation and performance*. https://doi.org/10.1007/978-3-030-49663-0_51
- Gras, B., Brun, A., & Boyer, A. (2020). *For and by student dashboards design to address dropout*. <https://hal.inria.fr/hal-02974682>
- Guzmán-Valenzuela, C., Gómez-González, C., Rojas-Murphy Tagle, A., & Lorca-Vyhmeister, A. (2021). Learning analytics in higher education: a preponderance of analytics but very little learning? *International Journal of Educational Technology in Higher Education*. <https://doi.org/10.1186/s41239-021-00258-x>
- Han, J., Kim, K. H., Rhee, W., & Cho, Y. H. (2021). Learning analytics dashboards for adaptive support in face-to-face collaborative argumentation. *Computers & Education*. <https://doi.org/10.1016/j.compedu.2020.104041>
- He, H., Dong, B., Zheng, Q., & Li, G. (2019, May 9). VUC. *Proceedings of the ACM Conference on Global Computing Education*. <https://doi.org/10.1145/3300115.3309514>
- Karaoglan Yilmaz, F. G., & Yilmaz, R. (2020). Learning analytics as a metacognitive tool to influence learner transactional distance and motivation in online learning environments. *Innovations in Education and Teaching International*. <https://doi.org/10.1080/14703297.2020.1794928>
- Kia, F. S., Teasley, S. D., Hatala, M., Karabenick, S. A., & Kay, M. (2020). How patterns of students dashboard use are related to their achievement and self-regulatory engagement. *Proceedings of the Tenth International Conference on Learning Analytics & Knowledge*. <https://doi.org/10.1145/3375462.3375472>
- Kitto, K., & Knight, S. (2019). Practical ethics for building learning analytics. *British Journal of Educational Technology*. <https://doi.org/10.1111/bjet.12868>
- Klerkx, J., Verbert, K., & Duval, E. (2017). Learning Analytics Dashboards. In *Handbook of Learning Analytics* (pp. 143–150). Society for Learning Analytics Research (SoLAR). <https://doi.org/10.18608/hla17.012>
- Knight, S., Gibson, A., & Shibani, A. (2020). Implementing learning analytics for learning impact: Taking tools to task. *Internet and Higher Education*, 45, 100729. <https://doi.org/10.1016/j.iheduc.2020.100729>
- Kokoç, M., & Altun, A. (2021). Effects of learner interaction with learning dashboards on academic performance in an e-learning environment. *Behaviour & Information Technology*. <https://doi.org/10.1080/0144929X.2019.1680731>
- Leitner, P., Maier, K., & Ebner, M. (2020). *Web analytics as extension for a learning analytics dashboard of a massive open online platform*. https://doi.org/10.1007/978-3-030-47392-1_19
- Leitner, P., Ebner, M., & Ebner, M. (2019). Learning analytics challenges to overcome in higher education institutions. In D. Ifenthaler, D. K. Mah, & J. Y. Yau (Eds.), *Utilizing learning analytics to support study success* (pp. 91–104). Springer International Publishing. [10.1007/978-3-319-64792-0_6](https://doi.org/10.1007/978-3-319-64792-0_6)
- Liu, R., & Koedinger, K. R. (2017). Going beyond better data prediction to create explanatory models of educational data. *The Handbook of Learning Analytics*. <https://doi.org/10.18608/hla17.006>
- Lu, J., Liu, A., Dong, F., Gu, F., Gama, J., & Zhang, G. (2018). Learning under concept drift: A review. *IEEE Transactions on Knowledge and Data Engineering*. <https://doi.org/10.1109/TKDE.2018.2876857>
- Mahroeian, H., Daniel, B., & Butson, R. (2017). The perceptions of the meaning and value of analytics in New Zealand higher education institutions. *International Journal of Educational Technology in Higher Education*. <https://doi.org/10.1186/s41239-017-0073-y>
- Maisarah, N., Khuzairi, S., & Cob, Z. C. (2020). The divergence of learning analytics research. *International Journal of Advanced Science and Technology* 29(6s). <http://sersc.org/journals/index.php/IJAST/article/view/9300>
- Majumdar, R., Akçapınar, A., Akçapınar, G., Flanagan, B., & Ogata, H. (2019). LAVIEW: Learning Analytics Dashboard Towards Evidence-based Education. In: *In Companion Proceedings of the 9th International Conference on Learning Analytics and Knowledge* (pp. 1–6). Society for Learning Analytics Research (SoLAR).

- Menon, S., & Suresh, M. (2020). Factors influencing organizational agility in higher education. *Benchmarking: an International Journal*, 28(1), 307–332. <https://doi.org/10.1108/BIJ-04-2020-0151>
- Moher, D., Liberati, A., Tetzlaff, J., & Altman, D. G. (2009). Preferred reporting items for systematic reviews and meta-analyses: The PRISMA statement. *BMJ (online)*. <https://doi.org/10.1136/bmj.b2535>
- Mothilal, R. K., Sharma, A., & Tan, C. (2020). Explaining machine learning classifiers through diverse counterfactual explanations. *FAT* 2020—Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 607–617. <https://doi.org/10.1145/3351095.3372850>
- Muhammad, R. N., Tasmin, R., & Nor Aziati, A. H. (2020). Sustainable competitive advantage of big data analytics in higher education sector: An Overview. *Journal of Physics: Conference Series*. <https://doi.org/10.1088/1742-6596/1529/4/042100>
- Mukerjee, S. (2014a). Agility: A crucial capability for universities in times of disruptive change and innovation. *Australian Universities' Review*, 56, 56–60.
- Mukerjee, S. (2014b). *Organizational agility in Universities* (pp. 15–25). IGI Global.
- Namoun, A., & Alshantiri, A. (2020). Predicting student performance using data mining and learning analytics techniques: A systematic literature review. *Applied Sciences*. <https://doi.org/10.3390/app11010237>
- Naranjo, D. M., Prieto, J. R., Moltó, G., & Calatrava, A. (2019). A visual dashboard to track learning analytics for educational cloud computing. *Sensors*. <https://doi.org/10.3390/s19132952>
- Nguyen, A., Gardner, L., & Sheridan, D. P. (2020). Data analytics in higher education: An integrated view. *Journal of Information Systems Education*, 31(1), 61–71.
- Owatari, T., Shimada, A., Minematsu, T., Hori, M., & Taniguchi, R. (2020). Real-time learning analytics dashboard for students in online classes. *International Conference on Teaching, Assessment, and Learning for Engineering (TALE)*. <https://doi.org/10.1109/TALE48869.2020.9368340>
- Rets, I., Herodotou, C., Bayer, V., Hlosta, M., & Rienties, B. (2021). Exploring critical factors of the perceived usefulness of a learning analytics dashboard for distance university students. *International Journal of Educational Technology in Higher Education*. <https://doi.org/10.1186/s41239-021-00284-9>
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2018). Anchors: high-precision model-agnostic explanations. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, 1527–1535. www.aaai.org
- Rubel, A., & Jones, K. M. L. (2016). Student privacy in learning analytics: An information ethics perspective. *The Information Society*, 32(2), 143–159. <https://doi.org/10.1080/01972243.2016.1130502>
- Sedrakyan, G., Malmberg, J., Verbert, K., Järvelä, S., & Kirschner, P. A. (2020). Linking learning behavior analytics and learning science concepts: Designing a learning analytics dashboard for feedback to support learning regulation. *Computers in Human Behavior*. <https://doi.org/10.1016/j.chb.2018.05.004>
- Sun, K., Mhaidli, A. H., Watel, S., Brooks, C. A., & Schaub, F. (2019). It's My Data! Tensions among stakeholders of a learning analytics dashboard. *Proceedings of the CHI 2019 Conference on Human Factors in Computing Systems*. <https://doi.org/10.1145/3290605.3300824>
- Tufte, E. *The Visual Display of Quantitative Information*, (2001).
- Ulfa, S., Fattawi, I., Surahman, E., & Yusuke, H. (2019). Investigating learners' perception of learning analytics dashboard to improve learning interaction in online learning system. *2019 5th International Conference on Education and Technology (ICET)*. <https://doi.org/10.1109/ICET48172.2019.8987229>
- Umer, R., Susnjak, T., Mathrani, A., & Suriadi, L. (2021). Current stance on predictive analytics in higher education: Opportunities, challenges and future directions. *Interactive Learning Environments*. <https://doi.org/10.1080/10494820.2021.1933542>
- Valle, N., Antonenko, P., Valle, D., Sommer, M., Huggins-Manley, A. C., Dawson, K., Kim, D., & Baiser, B. (2021). Predict or describe? How learning analytics dashboard design influences motivation and statistics anxiety in an online statistics course. *Educational Technology Research and Development*. <https://doi.org/10.1007/s11423-021-09998-z>
- Wachter, S., Mittelstadt, B., & Russell, C. (2017). Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *SSRN J*, 31, 841.
- Wilson, A., Watson, C., Thompson, T. L., Drew, V., & Doyle, S. (2017). Learning analytics: Challenges and limitations. *Teaching in Higher Education*, 22(8), 991–1007. <https://doi.org/10.1080/13562517.2017.1332026>
- Yoo, M., & Jin, S.-H. (2020). International forum of educational technology & society development and evaluation of learning analytics dashboards to support online discussion activities. *Technology & Society*, 23(2), 1–18. <https://doi.org/10.2307/26921130>

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Chapter 8

Effectiveness of a Learning Analytics

Dashboard for Increasing Student Engagement Levels

STATEMENT OF CONTRIBUTION DOCTORATE WITH PUBLICATIONS/MANUSCRIPTS

We, the candidate and the candidate's Primary Supervisor, certify that all co-authors have consented to their work being included in the thesis and they have accepted the candidate's contribution as indicated below in the *Statement of Originality*.

Name of candidate:	Gomathy Suganya Ramaswami
Name/title of Primary Supervisor:	Dr. Teo Susnjak
In which chapter is the manuscript /published work:	Chapter 8
Please select one of the following three options:	
<input checked="" type="radio"/> The manuscript/published work is published or in press <ul style="list-style-type: none"> • Please provide the full reference of the Research Output: Ramaswami, G., Susnjak, T., & Mathrani, A. (2023). Effectiveness of a Learning Analytics Dashboard for Increasing Student Engagement Levels. <i>Journal of Learning Analytics</i>. 	
<input type="radio"/> The manuscript is currently under review for publication – please indicate: <ul style="list-style-type: none"> • The name of the journal: • The percentage of the manuscript/published work that was contributed by the candidate: • Describe the contribution that the candidate has made to the manuscript/published work: 	
<input type="radio"/> It is intended that the manuscript will be published, but it has not yet been submitted to a journal	
Candidate's Signature:	Gomathy Ramaswami <small>Digitally signed by Gomathy Ramaswami Date: 2022.11.07 07:42:59 +13'00'</small>
Date:	07-Nov-2022
Primary Supervisor's Signature:	Teo Susnjak <small>Digitally signed by Teo Susnjak Date: 2022.11.07 15:30:10 +13'00'</small>
Date:	7-Nov-2022

This form should appear at the end of each thesis chapter/section/appendix submitted as a manuscript/publication or collected as an appendix at the end of the thesis.

Effectiveness of a Learning Analytics Dashboard for Increasing Student Engagement Levels

Gomathy Ramaswami^{1*}, Teo Susnjak^{2*}, Anuradha Mathrani^{3*}

Abstract

Learning Analytics Dashboards (LADs) are gaining popularity as a platform for providing students with insights into their learning behavioral patterns in online environments. Existing LAD studies are mainly centered on displaying students' online behaviors with simplistic descriptive insights. Only a few studies have integrated predictive components, while none possess the ability to explain how the predictive models work and how they have arrived at specific conclusions for a given student. A further gap exists within existing LADs with respect to prescriptive analytics which generates data-driven feedback to students on how to adjust their learning behavior. The LAD in this study attempts to address this gap and integrates a full spectrum of current analytics technologies for sense-making while anchoring them within educational theoretical frameworks. This study's LAD (SensEnablr) was evaluated for its effectiveness to impact learning in a student cohort at a tertiary institution. The study's findings demonstrate that the students' engagement with learning technologies and course resources increased significantly immediately following interactions with the dashboard. Meanwhile, the study results showed that the dashboard boosted the respondents' learning motivation levels and that the novel analytics insights drawn from predictive and prescriptive analytics were beneficial to their learning. This study, therefore, has implications for future research when investigating student outcomes and optimizing student learning using LAD technologies.

Notes for Practice

- Learning Analytics Dashboards (LADs) hold promise for enhancing the student learning environment by providing data-driven monitoring of students' learning progression and identifying overall learning trends and patterns.
- It is worthwhile integrating predictive analytics into LADs provided that transparency into how the models work and how they arrive at predictions for a given student is clarified.
- Data-driven suggestions offered by prescriptive analytics are appreciated by students and should be further researched and integrated into future LAD studies.
- Information-rich dashboards can incorporate multiple reference frames like the ability for students to compare themselves with their peers and monitor their levels of achievement towards their learning goals or progression with respect to their earlier selves without creating an overwhelming LAD which induces cognitive bias.

Keywords

Learning analytics, dashboard, descriptive analytics, predictive and prescriptive analytics, dashboard evaluation, student engagement behaviors, usability

Submitted: 05/12/22 — **Accepted:** — **Published:**

¹ ORCID ID: <https://orcid.org/0000-0002-5400-3843>, Email: g.ramaswami@massey.ac.nz

Corresponding author ² Email: t.susnjak@massey.ac.nz, ORCID ID: <https://orcid.org/0000-0001-9416-1435>

³ ORCID ID: <https://orcid.org/0000-0002-9124-2536>, Email: a.s.mathrani@massey.ac.nz

* Massey University, Auckland, New Zealand

1. Introduction

Current advancements in technology-enabled learning have provided educators with new opportunities for enhancing student learning by leveraging Learning Analytics (LA). LA provides a research lens for uncovering student learning patterns from an ever-increasing amount of student data to better understand the overall learning process and to more effectively support students

achieve their learning goals. LA Dashboards (LADs) are visual displays of information extracted from online learning environments that are arranged in such a way that students' learning trends can be perceived at a glance (Schwendimann et al., 2017). LADs, therefore, hold much promise for enhancing the learning environment by providing data-driven monitoring of students' learning progression so that actionable insights resulting in interventions can be taken to assist student learning.

Published LADs vary greatly in their complexity and sophistication of depicted insights. LADs based on *descriptive* analytics are the most common. They employ the simplest type of analytics that primarily depict current and historical data to identify trends. LADs that incorporate *predictive* analytics frequently leverage machine learning (ML) algorithms to provide forecasts about the likelihood of students achieving some predefined goal such as course or assessment completions (Shahiri & Husain, 2015; Rets et al., 2021). Meanwhile, *prescriptive* analytics capabilities produce data-driven feedback to students as well as remedial suggestions when necessary. LADs that incorporate prescriptive analytics are rare as this is the emerging frontier of current AI research (Susnjak, 2022).

The use of LADs is a relatively new adoption in education; it initially targeted educators to highlight students considered at risk of academic non-achievement. Purdue's Course Signals (Arnold & Pistilli, 2012) was one of the first widely deployed educator-facing dashboards that classified students into risk categories with respect to forecasted course outcomes. More recently, student-facing LADs have gained increasing interest in LA research (Teasley, 2017). LADs can facilitate more transparency into online learning behavioral patterns which can bring about new self-knowledge among students enabling them to make more informed study-related decisions for improving their learning outcomes (Verbert et al., 2013). However, distilling large amounts of multidimensional data into relevant information which can then be converted into meaningful and actionable insights for students is challenging (Khosravi et al., 2021). These challenges relate to the requirement that the LAD is both informative and comprehensive. To be informative, the LAD should be easy to understand and effectively convey insights that are consequential and judiciously selected and are not merely interesting. Whereas, to achieve comprehensiveness, a holistic overview of the student learning experience (Ali et al., 2013) needs to be depicted which combines data on learning content, learning activities, learning outcomes, and students' abilities in an effective manner that does not result in cognitive overload. Finding the right balance between informative and comprehensive visual displays is not a straightforward balancing act.

The effectiveness and proliferation of LADs in real-life use have also faced several challenges. Both educator- and student-facing LADs have increasingly incorporated predictive analytics to automatically identify at-risk students (Hu et al., 2014; Ahadi et al., 2015). However, this development has started to raise concerns centering on the prospect of widespread automated decision-making taking place with variable levels of human oversight (Hajian et al., 2016) while using machine learning (ML) models that are beyond obvious interpretation. Currently, published LADs that leverage predictive analytics do not provide users with any interpretation of their mechanics or any reasoning behind their predictions. Consequently, the development of accountable, transparent, and fair Artificial Intelligence (AI) technologies is being recognized as an important line of research in LA. Researchers in this field are interested in exposing the internals of black-box models to communicate the mechanics and provide explanations of the models' conclusions (Islam et al., 2022).

Additionally, while the sophistication and capabilities of the LADs have been increasing, the uptake and adoption of this tool have not matched the technological advances. It has been argued that one crucial factor determining whether students will use a LAD is their perception that the technology is predictable, reliable, and useful (Rienties et al., 2018). LAD research studies have posited that students will only engage with a LAD if they trust the data and understand how the predictions regarding them are formulated (de Quincey et al., 2019). Indeed, detailing the reasons why students receive a particular recommendation has been shown to increase their trust in the system and the likelihood of them following the advice (Bodily & Verbert, 2017). Although, when a model produces unexpected or erroneous output, trust is further eroded and could result in increased skepticism and possibly a rejection from the end user (Susnjak et al., 2022). These errors may have negative side effects as well, such as in instances when certain actions or decisions affecting others are taken, which might be based on false premises arising from misclassifications (Valle et al., 2021). While errors in predictions are not uncommon, they can to some degree be mitigated with explanatory technologies that unpack model behaviors. While crucially important, these kinds of explanatory capabilities are largely absent from present LAD implementations (Afzaal et al., 2021).

Lastly, a key challenge for LADs to improve their effectiveness has been in converting information into concrete and actionable steps for students. A central limitation of previous studies has been in developing prescriptive analytics capabilities that are embedded within LADs, which provide students with tailored suggestions on how to maximize their learning outcomes (Lepeniotti et al., 2020). Increasingly new tools are now becoming available from the eXplainable AI field which supports the next step in the evolution of LADs towards the incorporation of prescriptive capabilities that arguably embody the most potential for affecting improved learning outcomes of at-risk students.

This study seeks to contribute towards overcoming the above challenges and limitations of existing LADs by showcasing a student-facing LAD that integrates all levels of analytics as well as data-driven prescriptive analytics which offers tailored advice to students. Transparency of the underlying predictive model working is brought about through explanatory capabilities

that offer detailed explanations of how the model has arrived at specific predictions for a given student. Additionally, these advanced analytics capabilities are combined with comprehensive visualizations. Studies suggest that students' learning performance generally increases as their engagement level with the Learning Management Systems (LMS) also increases (Yu & Jo, 2014; Zacharis, 2015; Lu & Cutumisu, 2022). Therefore, this study aims to use measurements of engagement levels with LMS as a proxy of the overall learning performance. To that end, the goal of this article is to quantify the effect that this study's LAD has on students' engagement levels with the LMS, and in parallel qualitatively analyze students' subjective perception of the dashboard's effect on their learning performances.

Schwendimann et al.'s (2016) review of LADs emphasizes the importance of grounding their designs in established educational theoretical frameworks. Jivet et al. (2018) note that only a few reviewed papers outlining LADs have taken account of theories within their design, while Verbert et al., (2020) call for more responsible LA designs with pedagogical underpinnings in the development of LADs. To that end, a further key contribution of this paper is the rich grounding of LAD development within existing educational theory frameworks. Social Comparison Theory and Self-regulated Learning, in particular, have informed the design and development of this study's LAD, SensEnablr, which enables learners through sense-making to take actionable steps. By situating this study's LAD within these pedagogical contexts and others, the study seeks to bridge the gap between LADs in general and educational theory in order to promote the use of this technology to enhance learning.

The study addresses the following research questions

1. Does the SensEnablr analytical dashboard result in more student engagement with the learning materials within the institutional LMS?
2. What are the students' perceptions of the dashboard's effect on their learning performance?

2. Background

This section lays out the current state of research in LADs with an investigation into the different levels of analytics used by LADs in the literature. The literature review also considers the breadth of evaluations conducted on the effectiveness of published LADs.

2.1 Learning analytics dashboards

In recent years, LADs have rapidly evolved and have augmented simplistic descriptive capabilities by incorporating predictive analytics components. Hellings & Haelermans (2020) designed a dashboard aiming to provide study progress updates to students along with their predicted probability of success in a course accompanied by the predicted grade. Linear models were used to predict the grade mark, while AdaBoost predicted the course outcomes. A weekly email with the dashboard link was sent to the participants with the intention of encouraging them to use the dashboard frequently. Dashboard usage was correlated with a positive impact on student online engagement, but no impact on final exam grades and course completions was detected. Similarly, Baneres et al. (2019) focused on devising an early warning system for students and educators that identified at-risk students. Their Graduate At-Risk (GAR) model used grades to predict course-level outcomes. GAR noted an improvement in the performances of the at-risk students; however, the effects could not reliably be attributed to the tool itself as the primary factor or the interventions.

In a separate study, a descriptive and predictive analytics dashboard was developed by Valle et al. (2021). The descriptive dashboard displayed the students' performance relative to the class average while the predictive dashboard displayed the probability of students attaining specific grades. The authors reported that the predictive dashboard positively impacted only the highly motivated students to sustain their motivation levels, although both dashboards failed to demonstrate their effectiveness in affecting final outcomes. Recently, Duan et al., (2022) designed a LAD to provide students with actionable feedback on their weekly learning progress to advance their self-regulated learning skills and improve their course performance. A detailed inquiry using mixed methods was also conducted to study the dashboard's impacts on students. It was found that students' use of the dashboard was positively correlated with their course performance, and those who viewed the dashboard had higher course ranks. In addition, the positive correlation associated with the use of the dashboard also extended to more timely homework submissions.

However, while such descriptive- and predictive-based dashboards provide interesting and potentially revealing information, which is typical of most existing LADs, importantly they do not provide specific guidance and recommendations (such as guiding students towards relevant learning materials or activities that are likely to increase course performances) nor do they provide actual explanations of the predictions which may hold clues regarding possible remedial actions that students can take.

2.2 Predictive model explainability

Jayaprakash et al. (2014) argue that awareness of predictive models' outputs alone does not influence course completion and

retention rates unless these are combined with effective intervention strategies aimed at supporting at-risk students. However, it is also widely appreciated that predictive models have the potential to provide timely intervention for at-risk students which can result in corrective measures being undertaken by them (Namoun & Alshantqi, 2020). Extracting maximal value from predictive models is the goal. Recent advances in machine learning technologies have shown how this can be achieved by not only enabling explanations of models and their predictions but also by generating data-driven and automated counterfactuals which can be converted into tailored prescriptive feedback. Ultimately, these outputs can articulate to the students what behavioral adjustments would hypothetically result in positive predictive outcomes. Such explanations hold potential in helping students regulate their online behavior in a data-driven manner (Afzaal et al., 2021). Thus, arguably the most beneficial and insight-rich form of analytics is found in the prescriptive data-driven outputs which generate the greatest intelligence and value (Lepenioti, 2020).

As it currently stands, the focus of existing LADs (Baneres et al., 2019; Hellings & Haelermans., 2020; Valle et al., 2021) has been merely on conveying to students their at-risk status or probability of non-completion in a course. None of the existing predictive LADs adequately explain to users how the models work, or how their predictions were generated; rather, they merely provide the predicted outcomes as feedback to users. This compounds the challenges facing the acceptance and uptake of LAD technology since it has been noted that students would use these types of tools more frequently if they understood the outputs better and if they were provided with clear evidence that they are an effective tool (Kim et al., 2016).

Feedback is only effective if students can act on it (Ryan & Henderson, 2018). Several studies (Baneres et al., 2019; Kokoç & Altun, 2021; Li & Fan, 2019) have leveraged the dispatch of feedback in the form of manually generated messages and recommendations directly from educators. However, these examples of feedback were not individually tailored and did not specifically identify to the students the learning strategies they ought to employ as a remedial response. Algorithmic approaches to automatically generate evidence-based and data-driven prescriptive feedback, which is personalized, therefore, hold great promise for improving outcomes.

2.3 LAD impact

Examining the impact and perceived usefulness of LADs on student online behavior, achievement, and skills is considered important (Bodily & Verbert, 2017) not least because these tools represent significant resource investments. But naturally also, because there is a need to quantify the effect of these tools to mitigate general worldwide trends in decreasing rates of learning performances and retention in tertiary education.

Few studies (Baneres et al., 2020; Fleur et al., 2020; Kokoç & Altun, 2021) which focused on generating the student risk status to identify the at-risk students have claimed that usage of the LAD showed a positive impact on students' outcomes to those students who used the dashboards. Other studies (Bodily et al., 2018; Chatti et al., 2020; Han et al., 2021) have used qualitative approaches to evaluate general usability aspects like assessing users' perceptions of the tool based on their interactions. Bodily and Verbert conducted a review on a student-facing LAD and concluded that student use of LADs is generally not well studied nor adequately understood, and not having enough student evaluations "is detrimental to the research field of learning analytics because a lack of usability could be the reason why students do not like or use a system" (p. 417). The lack of student evaluations restricts us from knowing whether current LAD systems are adequately meeting the goals for which they have been designed. While recent studies like that of Duan et al. (2022) provided insights into the LAD's impact on course performance and homework submission time, more comprehensive evaluations would help researchers understand the full extent of the LAD's effect on self-regulation and learning strategies. Thorough student evaluations can therefore highlight end-users' perceptions regarding LAD usability, its usefulness, and the impact of the LAD on self-regulating their learning strategies. Moreover, evaluations can help educators understand whether LAD recommendations are being properly communicated to students and whether those students feel motivated by the offered guidance.

3. Theoretical Framework

The design of the LAD in this study is supported and situated within various theoretical frameworks in education. Schwendimann et al. (2016) stressed the importance of anchoring LAD design in established educational theoretical frameworks. Their systematic literature review on LAD research identified key design features and explored the extent to which they were linked within established educational theoretical frameworks. Their work found a very few LAD designs and implementations had articulated how their work was grounded within existing theoretical frameworks. Overall, they identified six broad types of indicators related to learners, actions, content, context, result and social, although how these indicators were linked to theory was not evident in the reviewed studies. Guiding LAD designs with educational frameworks, such as Self-regulated Learning (SRL), Social Comparison Theory (SCT), and what could be termed as Cognitive Load Theory (CLT), or

Feedback Intervention Theory (FIT) can demonstrate the benefits of learning and teaching. The review in this study shows that LADs grounded in SRL principles are more common, since these aim to provide learners with performance, progress, and strategy information to facilitate self-reflection and learning behavior adjustments. A limited number of studies incorporated elements from the Social Comparison Theory (SCT), displaying student data in relation to that of anonymized peer performance data to trigger greater motivation. Others still were designed with CLT principles, focusing on clear visualizations and actionable feedback to prevent information overload. Finally, a small number of studies referenced FIT, with dashboards providing real-time, personalized feedback to support learners in identifying areas of improvement and enhancing their performance.

The design of SensEnablr draws heavily from the SCT and SRL theories, but is also grounded within Constructivism, SCT, and Transformative Learning (TL), which is outlined below. These theories help us understand how some of the more novel features of this study's LADs can be more rigorously justified and used effectively to enhance students' self-awareness, self-regulated learning, and raise academic achievement.

3.1 Social Comparison Theory

The intention of LADs is to enhance students' self-awareness, leading to improved self-regulated learning, and ultimately, enhanced academic achievement (Lim et al., 2019). From a theoretical standpoint, for LADs to be effective, students require some form of "representative reference frame", for interpreting their data on a dashboard (Wise, 2014). The term "frame of reference" is an internal/external frame of reference (also known as the I/E model) proposed initially by Marsh (1986). The model proposes that academic self-conceptualization or students' perception of their own academic abilities is shaped by internal and external comparisons. Internal comparison relates to measuring a student's achievements across dissimilar learning domains (i.e., different courses), while external comparison measures a student's achievement with that of their peers in the same learning domain. By using reference frames, students can build a self-concept or self-perception of their academic competence, as they gain awareness of their own strengths and weaknesses, which often leads them to refine their learning profiles (Skaalvik & Skaalvik, 2004). Academic self-concept is found to be highly related to achievement, interest and aspirations, as increased perceptions of competency lead to increased levels of intrinsic motivation and vice versa, therefore, serving both as cause and effect (Marsh et al., 2005). Therefore, an effective communication strategy that mutually reinforces academic self-concept and academic achievement among students is crucial.

Current LADs too aim at increasing students' motivation by presenting them with dimensional achievement comparisons that reflect both internal and external reference frames. Jivet et al. (2017) categorized student-facing LADs according to three types of reference frames, namely (i) social (comparison with peers), (ii) achievement (distance towards goals), and (iii) progress (comparison with an earlier self). The social comparison can be made with the whole class, top students, and peers with similar goals. The achievement theory allows students to compare their performance level with desired goals from an internal perspective (i.e., set by themselves) or external perspective (i.e., set by their teacher). This frame might suit both performance goal-oriented students and mastery goal-oriented students because this frame makes it possible to focus on the learning outcomes and how they are mastering the materials and tasks. Progress allows learners to compare historical data of their performances with their current level of performance (Pintrich, P. R. (2000)).

The first is expressed as an external frame of reference, while the second and third can be interpreted as internal frames of reference, but from a slightly different angle by using progression over time rather than as direct comparisons between various learning domains or courses. Social comparison is the most widely used reference frame for LADs as this can reveal learning behaviors among peers with similar goals and knowledge.

However, few studies have analyzed the effects of using external frames of reference in LADs on learning outcomes. Jivet et al. (2018) note mixed results in using comparisons with peers between low and high-performing students. For low-performing students, the social comparison frame was perceived as being stressful. In contrast, the high-performing students were motivated by seeing their success in comparison to others. Differences in the way students respond to frames of reference may be related to baseline motivation or self-perception of one's competency. In another study, Roberts et al. (2017) noted that while students favor the utilization of an external frame of reference which enables them to compare their progress with peers for self-evaluation, the way this information is presented needs to be carefully considered by being cognizant of the fact that for some students this could be detrimental.

3.2 Self-regulated Learning Framework

Self-regulated learning (SRL) is a theoretical framework that emphasizes the role of learners as active agents in managing and monitoring their learning processes (Zimmerman, 2000). SRL involves setting goals, planning, organizing, monitoring progress, and evaluating outcomes to optimize learning and performance (Pintrich, 2000). Within the context of LADs, the SRL

framework can be a guiding principle in design and functionality to support learners in regulating their efforts. That is, LADs informed by SRL principles can typically provide learners with automated feedback about their performance, progress, and learning strategies, enabling them to reflect on their learning behaviors and adjust their approaches as needed (Howell et al., 2018). By offering actionable insights and fostering self-awareness, LADs grounded in SRL theory can empower learners to take greater control over their learning process, ultimately enhancing their motivation, engagement, and overall academic performance (Winne & Hadwin, 1998).

3.3 Further theoretical frameworks

Besides SCT and SLR, which heavily influenced the design of SensEnablr, other theories are also relevant. Constructivism posits that learners actively construct their own knowledge by building on their prior experiences and actively engaging with new information (Piaget, 1952; Vygotsky, 1978). Incorporating constructivist principles into LAD design involves providing students with meaningful and relevant information about their learning progress, which enables them to make sense of their performance, identify gaps in their understanding, and adjust their learning strategies accordingly. To achieve this, LADs can present data in a visually appealing and easy-to-understand manner, allowing learners to analyze their performance from multiple perspectives and recognize patterns in their learning behaviors. This process fosters self-reflection, critical thinking, and knowledge construction, which are essential elements of constructivist learning.

Moreover, LADs can be designed to promote social interaction and collaborative learning, aligning with Vygotsky's (1978) emphasis on the importance of social processes in knowledge construction. For example, dashboards can provide features that enable students to share their insights, experiences, and progress with their peers or engage in discussions about their learning. This exchange of ideas and experiences fosters a constructive learning environment, where learners can benefit from one another's perspectives and co-construct knowledge (Stahl et al., 2006).

In addition to individual and collaborative learning, LADs designed within the constructivist framework can also support the role of educators in facilitating the learning process. By providing teachers with insights into students' performance, engagement, and learning strategies, LADs can help educators identify areas where learners need additional support, enabling them to intervene and provide guidance that is tailored to the specific needs of each student.

On the other hand, Social Cognitive Theory (SCT) emphasizes the reciprocal relationship between personal factors, environmental influences, and behavior (Bandura, 1986). According to this theory, self-efficacy or one's belief in their ability to achieve a goal plays a critical role in motivation, learning, and achievement (Zimmerman, 2000). Firstly, LADs can foster self-efficacy by providing students with feedback on their performance, allowing them to monitor their progress, set realistic goals, and develop effective learning strategies (Zimmerman & Schunk, 2001). By presenting data on their accomplishments and areas for improvement, LADs can help learners gain a better understanding of their capabilities and adjust their efforts accordingly. This targeted feedback can enhance motivation, as students recognize their potential for success and become more confident in their ability to achieve their goals.

Moreover, LADs can incorporate social comparison features to promote vicarious learning, another key element of SCT. Vicarious learning occurs when students observe and learn from the successes and challenges of their peers (Bandura, 1986). LADs can display anonymized peer performance data, enabling learners to compare their progress and performance with that of their classmates. This comparative information can motivate students to improve, as they strive to achieve similar levels of success or avoid the pitfalls experienced by their peers.

Lastly, Transformative Learning Theory (TLF) posits that learners can experience a profound shift in their perspectives, values, and beliefs through critical self-reflection and dialogue (Mezirow, 1997). LADs can facilitate transformative learning by incorporating various features, including predictive and prescriptive analytics capabilities.

Predictive analytics within LADs enable learners to anticipate potential challenges and future performance based on their historical data and patterns. By identifying trends and uncovering potential obstacles, students can engage in critical self-reflection, reassess their learning strategies, and make proactive decisions to address potential issues. This process helps learners challenge their existing assumptions about their learning capabilities, fostering transformative learning experiences (Kitchenham, 2008).

Prescriptive analytics capabilities in LADs can provide students with personalized recommendations and actionable insights based on their performance data. These insights might include suggestions for improving study habits, prioritizing specific course content, or seeking additional support. By presenting these tailored recommendations, LADs with these capabilities empower learners to reflect on their current approaches and consider alternative strategies, potentially leading to a shift in their perspectives and beliefs about their learning (Cranton, 2006).

4. Dashboard design

In light of the discussion of theoretical underpinnings connecting LADs with established theory, components of SensEnablr are presented here and linked to the reviewed frameworks.

The layout of SensEnablr can be seen in Figure 1, reflecting a comprehensive integration of both internal and external reference frames while the various components are anchored in the reviewed theoretical frameworks. The descriptive, predictive, and prescriptive components are structured into three columns. The first column shows what is primarily a descriptive component that leverages the external frame of reference, drawing heavily from SCT to enable a comparison with peers, which prior works have found to be rarely used in LADs (Jivet et al. (2018)). Meanwhile, the visualizations in this column also enable monitoring one’s progress over time thus drawing on SRL as well. In this panel of visualizations, a student’s engagement level is compared against the backdrop of their class’s average. The engagement value is a combination of the number of login counts into Moodle, the number of learning resources accessed, and the total forum posts created, measuring communication exchanges with peers. By providing the learner with a reference point to their peers, the intention is to enhance students’ self-awareness with the goal of leading to improved self-regulated learning, drawing also from the SCT.



Figure 1. The SensEnablr LAD used for this study

To highlight more clearly aspects of the charts in the first column, Figure 2, presents the weekly engagement scores of a specific student in comparison to the class average over the first seven weeks of the semester, illustrating the application of SCT principles. The engagement scores are derived from various indicators, including LMS login counts, forum discussion activities, and access to online learning materials, reflecting the student's level of interaction with the online learning environment. The student's engagement score consistently falls below the class average throughout the seven-week period, representing an external reference frame that the student can use for comparison. This trend suggests that the student may be less engaged with the course materials and online discussions compared to their peers. By leveraging SRL principles, the student can use this external reference frame to set re-adjusted goals, continue monitor their progress against the new goals, and continue to revise their learning strategies accordingly. Simultaneously, the student can employ an internal reference frame, focusing on their personal growth and improvement.




Figure 2. The enlarged chart of a student’s engagement levels from the SensEnablr LAD

The second column features both descriptive and predictive components; similar to the first column the descriptive components have followed the external frame of reference, while also considering the principles of Constructivism. It displays a snapshot of the student’s assignment grades and test marks within the context of the whole class. Figure 3 for example highlights one component from this section of SensEnablr where the student’s marks across two assessments are contrasted with those of the class average showing a positive trend for the student.

The real-time feedback and information about their learning performance enables learners to actively construct their own understanding of the knowledge and skills they have attained, and make adjustments to their learning strategies as needed. In this manner, learners are empowered to use the dashboard to identify areas where they may need additional support or to track their progress towards set learning goals.




Figure 3. The enlarged chart depicting a student’s assessment marks from the SensEnablr LAD

The dashboard’s predictive component begins in the bottom half of the second column. The students’ approximate upcoming assessment scores are presented along with estimated scores for their final exam. These estimations are derived based on how students with similar learning attributes and performances from the same course in the previous year had performed, akin to a k-Nearest Neighbor approach.

The predictive component that begins here, draws from the principles of TLF where learners can critically reflect on their own learning progress and identify areas where they may need to challenge their assumptions and beliefs when confronted with a prognosticated outcome that they may not expect.

Finally, the key predictive and prescriptive components are seen in the third column which continues both the application of the TLF as well as the internal reference frame (achievement theory). In this column, the student’s overall course outcome prediction is displayed. In addition, the underlying model accuracy is communicated along with a simplified version of model interpretability to the one described in a subsequent section in Figure 5, highlighting only the top features and their relative strengths when formulating predictions.

Furthermore, in the last column, the dashboard contains model explainability that attempts to describe to students how the model had arrived at a given prediction for their case. The internal reference frame (distance towards goals) is reflected in the lower half of the third column where the above model transparency capability is further extended with the use of counterfactuals that provide the dashboard’s prescriptive components. In this section, the dashboard offers recommendations to students on what modifications can be made to their learning behaviors for maximizing their learning outcomes. The automated feedback generated by counterfactuals offers personalized suggestions that are data-driven and can easily be configured to generate human-understandable text.

4.1 Dashboard architectural layers

In building the dashboard, the software artefact comprises three technological components, namely the Data Layer, the Data Analytics Layer, and the Presentation Layer. The high-level architecture with a description of the main tasks of each layer can be seen in Figure 4. Meanwhile, there are two separate datasets (detailed in Sections 5 and 6) that were processed by this architecture and served different purposes. The first dataset (modelling dataset) comprised past students whose data was used to develop the predictive and prescriptive models. The second dataset (dashboard dataset) comprised students participating in the LAD trial, which also included data from their class peers in an anonymized form.

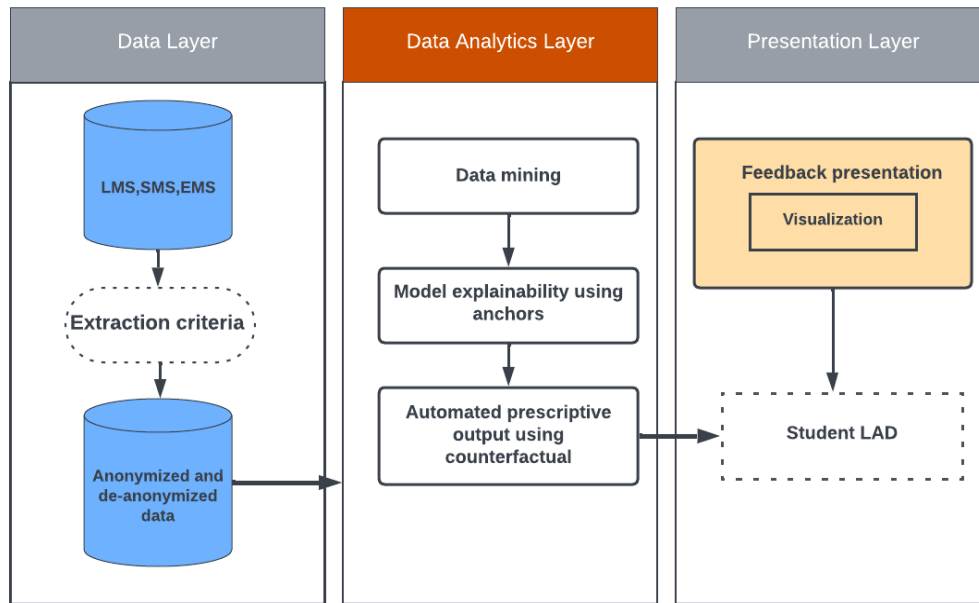


Figure 4. LAD implementation architecture

In the first component, the Data Layer, the raw data comprising both the modelling and the dashboard datasets are extracted from several databases, LMS (Moodle), Student Management System (SMS), and the Enrolment Management System (EMS), and filtered based on target years and courses. The extracted modelling dataset is anonymized in its entirety, while the dashboard dataset is anonymized only to shield the identities of students who have not opted-in to participate in this trial. The data is then cleaned and new features describing each student are engineered yielding a single dataset that is able to support the subsequent analytics tasks. The data extraction and processing step is conducted only once for the modelling dataset, while it is repeated on a weekly basis for the dashboard dataset in order to maintain an updated version of the LAD.

The Data Analytics Layer consists of three tasks that support the accountability and transparency goals of predictive modelling as well as those of prescriptive analytics. First, the Data mining step employs machine learning to build predictive models that identify at-risk students of course non-completion by using the modelling dataset. This represents a one-off task which does not need to be repeated. The resulting models in this step are applied to the dashboard data for participating students in order to generate their predictions as well as the explainability outputs that described how and why the model has produced given outputs. Lastly, techniques that enable the automated construction of prescriptive feedback to the participating students are used.

The third component is the visual aspect of the dashboard where the aim was to integrate all the descriptive, predictive and prescriptive components into a graphical representation. From a technical point of view, the visual aspect of SensEnablr was implemented following a client-server architecture using the Power BI web application tool (Knight et al., 2018) on the client side. The web application was populated with the data for each student and protected through authentication via logins. The dashboard was built on Power BI Desktop and subsequently deployed using the Power BI service that is built on Microsoft's Azure cloud computing infrastructure. Meanwhile, Python's scikit-learn, and other libraries were used for building all the data layer processing and analytics tasks. More detailed explanations of these processes are discussed in the next section.

4.2 Dashboard access

Only participating students who provided informed consent to opt into the study were given access to the dashboard which was advertised to a large body of students. The consenting students received the dashboard link via email, which allowed them to view their weekly performance. A manual guide explaining the login process and the usage of SensEnablr was also provided.

Students who did not consent to participate in the study did not have access to any parts of the LAD. To maintain privacy and confidentiality in compliance with the institutional Ethics Committee instructions, teachers and administrators were not granted access to the dashboard either. Only the research team which processed the data and developed SensEnablr, and the opted-in students could view the dashboard, while those students could only see information specific to themselves. SensEnablr was made available to the participants four weeks into the semester and remained accessible until the end of the semester.

4.3 Dashboard analytics

The nature of the predictive model used in the dashboard is outlined here. The focus is on describing the features (or variables) used for communicating the accuracy of the models achieved across several machine learning algorithms. The internal mechanics of the chosen model are exposed with respect to the impacts that the selected features have on the predictions and thus model interpretability is introduced. This section shows how this type of analysis is translated to the dashboard and it also discusses the technologies used to extract the explainability of the models to make the reasoning of the predictions transparent to the students. Finally, methods used to embed prescriptive analytics aspects into the dashboard and the approach to automate this functionality are described.

4.3.1 Predictive analytics

The training of the predictive models was conducted on anonymized data from the 4,000 historic students. New features were engineered to build the classification models. The aim was to engineer *course-agnostic* and thus generic features describing each student which were not tightly coupled to the particulars of different courses. Therefore, where possible the engineered features captured the relative values of a student compared to the mean values of their class. To achieve this, z-score (or standard score) was calculated on LMS data which relativized a student’s score on a given feature with respect to those of their peers, capturing the degree of deviation from the class mean. In general, feature values for each student were calculated on a rolling mean basis where actions performed by students over multiple weeks were used as a single column, hence reducing the number of feature columns. This facilitated a reduction in the total number of features which helped mitigate the model overfitting problem. All the used features are outlined and represented in Table 1.

Table 1. Feature description

Feature name	Description
Average score of prior courses	The mean score achieved by a student across all previous courses, reflecting their academic performance history
Maximum score achieved in prior course	The highest score achieved by a student in any of their previous courses, indicating their peak academic performance
Prior course deviation score	The z-score of a student's prior course scores, representing how much their performance deviates from the class mean
Assignment score	The accumulated assignment scores received by a student in the current course
Assignment deviation score	The z-score of the student's mean assignment score, indicating how much their assignment performance deviates from the class mean
Prior role description	Student's primary activity during the previous year with respect to their current academic year (e.g., not employed, beneficiary, self-employed, wage or salaried worker, secondary school student, polytechnic student)
LMS deviation score	The z-score of a student's engagement score, expressing how much their engagement deviates from the class mean where engagement is calculated based on login counts, forum discussion activities, access of online learning materials
LMS engagement score	The total count of activities performed by a student on the Moodle platform, reflecting their level of interaction with the LMS based on login counts, forum discussion activities, access of online learning materials as a rolling 4-week average
Citizenship	Student's nationality, which may influence their cultural and educational background
Age	Age of a student, potentially impacting their life experiences and learning habits
High school qualification	The highest school qualification a student obtained at admission, providing a baseline of their prior academic achievement
Study mode	Whether a student studies through distance/online learning or on-campus, which may affect their learning experience and engagement
Gender	Gender of the student, as it may play a role in their learning preferences and experiences
English proficiency test	Student's English language proficiency level, as determined by standardized tests (e.g., IELTS, PTE, NZCEL Level 4, TOEFL), which may influence their ability to engage with course materials and communicate effectively

The predictive problem for the primary model was formulated to predict a given student as being either "High" or "Low" risk category with respect to successfully completing their course. To achieve this, all the historic students in the dataset who gained 60% or less in their total course marks were labelled as high risk and the remainder as low risk. The models were then trained to map the features to the two outcome categories. Experiments were performed with a wide range of algorithms, namely CatBoost (Dorogush et al., 2018), Random Forest (Breiman., 2001), Naïve Bayes (Domingos & Pazzani, 1997), Logistic Regression (Hosmer & Lemeshow, 2000) and k-Nearest Neighbours (Hechenbichler & Schliep), from which the best performing algorithm was chosen for use in the dashboard.

A modified k-fold cross-validation approach was used to evaluate the models. The study’s dataset comprised seven different courses with a total of 10 separate deliveries across them. It was determined to train a model using nine course deliveries and to test each model against the remaining hold-out course offering. This process was repeated 10 times with different combinations of training and hold-out courses to arrive at final, aggregated evaluation scores. Measures used were total accuracy, the F-measure, and the Area Under the Curve (AUC) for comparisons. Table 2 lists the accuracies of all the models from best performing to the least accurate according to the F-measure scores which are more reliable on datasets with imbalanced class labels since this approach balances the precision and recall values. Given that CatBoost achieved the highest scores, the models from this algorithm were selected to identify the at-risk students and to display its outputs on the dashboard. While comparing the accuracies of models across different studies is not always strictly valid, a recent systematic review (Namoun & Alshantiti, 2020) has found that average accuracies currently range between 75% and 95%, thus validating that the generated models are within expected norms of reliability. The proposed prediction model developed using CatBoost was used to predict the future potential at-risk students.

Table 2. Performance scores of various classifiers together with the standard deviations.

Classifiers	F-measure	Accuracy %	AUC
CatBoost	.77±.02	78±2.1	.87±.02
k-Nearest Neighbors	.71±.02	71±2.4	.72±.02
Naïve Bayes	.68±.02	68±2.3	.71±.03
Logistic Regression	.67±.03	68±3	.73±.03
Random Forest	.67±.03	67±2.4	.74±.02

4.3.2 Model interpretability

The behavior of the model was validated by examining which features are most impactful and how their changing values influence the final prediction. For this purpose, the Shapely Additive Explanations (SHAP) method (Lundberg, 2019) is used, with its output depicted in Figure 5. A simplified version of this kind of output is presented to the students on the dashboard. In Figure 5, features are ranked and listed from the most impactful downward. The top four features are dominated by aspects of students’ performance in previous courses and their assignment scores in their current course which testifies to the reasonableness of the model.

Within Figure 5, there is a secondary dimension that yields further insights regarding the model’s mechanics corresponding to changing feature values. The y-axis color gradient depicted on the right-hand side represents feature values, with green indicating high and red low values. The vertical line centering on zero on the x-axis represents the neutral effects of each feature on the final prediction. As the points for each feature extend to the right of the vertical line, the stronger the effects of that feature on predicting low risk for a given student becomes, and conversely in the left direction. Combining this with the color gradient, it can be observed that as the average score of a student in their prior course increases, the higher the propensity for the model to predict a student as low risk becomes, and likewise for the opposite scenario.




Figure 5. Feature importance plot with SHAP

4.3.3 Model explainability

While the high-level view of the behavior of predictive models shown in Figure 5 is important, it does not have the ability to translate a prediction for a given student into a precise and clear explanation which is needed for the dashboard in this study. Though SHAP has the ability to explain individual predictions together with other technologies like LIME (Ribeiro et al., 2016), in this study it was decided to instead use Anchors (Ribeiro et al., 2021) which reduce a black-box prediction into a simplified rule-based description consisting of simple IF-ELSE statements. An example of this type of output is shown in Figure 1 with the presentation of the dashboard.

4.3.4 Prescriptive analytics

Prescriptive analytics was implemented via the use of counterfactual modelling. This means that for a student identified as high risk, an alternative outcome was modelled for their scenario. Using their current input values and the trained predictive model, a minimal set of changes to the student’s input values were modelled that would need to change in order for the predictive model to toggle its forecast to low risk. For this, counterfactual modelling was constrained to only the features which are actionable, thus ignoring immutable features like citizenship, etc. The outputs of the counterfactual modelling were then converted into textual output that takes the form of understandable feedback.

5. Methods

This section describes the dataset for training the models as well as the methods used for the recruitment of the study participants and subsequent evaluation of the SensEnablr dashboard.

5.1 Dataset

The dataset used for training both the predictive and prescriptive models in this study consisted of undergraduate and postgraduate students from an Australasian education provider between 2016 and 2019. As such, the data used for creating the models consisted of historic learners whose data was anonymized. The total size of the training dataset included 4,000 students across four years.

The student population in the dataset was diverse, covering three different disciplines: Business, Health, and Sciences. The gender distribution among the students was as follows: 65% female, 34% male, and 1% gender diverse. Additionally, the dataset included both full-time and part-time students, with approximately 54% of the students enrolled part-time and the remaining 46% as full-time students.

This comprehensive dataset provided a robust foundation for training the predictive and prescriptive models employed in SensEnablr. The diverse nature of the dataset allowed the models to capture a wide range of student behaviors and characteristics, ultimately contributing to the effectiveness and accuracy of SensEnablr's predictions and prescriptions.

5.2 Participants

Educational datasets used for building SensEnablr predictions followed institutional codes of conduct; hence an institutional-level research ethics protocol sets out boundaries on students' data usage and limits authorized access to the available datasets. It further puts forth measures on maintaining learner privacy and preventing disclosure of past and current student data using de-identification methods (Mathrani et al., 2021). Accordingly, ethics approval was sought from the Human Research Ethics Committee at the host HE institution, and subsequently, approval to proceed was obtained both for building the prediction model using anonymized and aggregated student data and to further conduct research with those students who had consented to participate. Current students were informed about the nature and purpose of the proposed dashboard and invited to participate. Study participants comprised only those students who had voluntarily opted for access to the dashboard over the course of their study at the tertiary institution.

Approximately ~500 students across five different courses (2 from Business, and 3 Sciences) were initially invited to participate in this study. The majority of the participants were male (74%) and the remaining 26% were female. The study was advertised to students in-person during the deliveries of lectures as well as through LMS forum posts. Students were invited to email the research team if they wished to participate. In total, 30 student participants who had given their informed consent were given access to the dashboard, while other students could not access the dashboard. Moreover, to maintain students' privacy and confidentiality, teachers too were not given access to the dashboard. Only the research team and the opted-in students could view the dashboard. Each participant was provided with the dashboard link through email to view their weekly performances. The participants were provided with a manual guide explaining the login process to access the dashboard. Nearly 59% of the participants were classified as "High-risk" in week 4 i.e. before giving access to the dashboard. Meanwhile, by week 6 of the semester this fell to 39%.

All participants were familiar with the institutional LMS since they had earlier enrolled in courses that utilized the LMS. The said institution makes extensive use of Moodle (an open-source LMS) over which course resources (including digital book chapters, videos, and lecture recordings) are made available weekly to the students. In addition, discussion forum activities and learning tasks are designed for enabling students to jointly reflect upon the course content. Some courses made use of weekly quizzes that served as assessment activities. Students' interactions with the LMS, such as their click-events as they access course resources were recorded automatically in the system's database, as were the number of logins to the learning dashboard.

Course duration typically lasts 16 weeks in a given semester. Four weeks following the commencement of the semester, dashboard access was enabled for participants and continued until the end of the semester. Over the semester, students were tasked with quizzes and assignments that were then graded and the marks allotted. Each student's final course marks are calculated from weighted averages of marks obtained in various quizzes and assignments as well as those from the final examination. While SensEnablr accessed and processed data from all students in a given course, only the opt-in students had access to SensEnablr, and they could only see information specific to them. Instructors also did not have access to SensEnablr.

5.3 Dashboard

Assessing the impact and usefulness of SensEnablr was a target for evaluation. Statistical approaches were used to evaluate the impact of the dashboard on student engagement with LMS to answer the first research question. A paired sample t-test also known as a dependent sample t-test was performed to find whether significant differences existed between the two related variables (Zimmerman, 1997). In this study's case, the test compared the online behavior of the participants with Moodle before and after using the dashboard. The purpose was to determine whether displaying the e-learning activities of the participants on the LAD had encouraged students to engage more with Moodle. Three different time frames of 15-day intervals were chosen for evaluation, each one designating the before and after conditions of using the dashboard. This analysis enabled us to make some inferences about the effects of the dashboard on short-term learning behavioral patterns.

5.4 Survey

In addition to an objective evaluation, a subjective evaluation of SensEnablr was also conducted. This involved assessing the students' perception of the dashboard's effect on their learning performance. 30 survey responses were collected. The system ability scale (SUS) has been used for evaluating the dashboard's overall usability (Brooke, (1996)) which is known to give a reliable measure of the perceived usability of a system even with a small sample (e.g., 8-12 users) (Tulis & Stetson, 2004; Brooke, 2013). The participants were asked to provide responses to the following questions:

1. I would like to continue to be able to use this dashboard for improving my performance

2. The dashboard presents relevant information that users can easily understand at a glance.
3. The dashboard is user-friendly and intuitive to use
4. Insights about my learning activities from the dashboard motivated me to continue learning
5. The dashboard made me aware of my current online engagement behavior and helped me improve
6. The prediction of my risk-profile had a positive impact on my attitude toward my studies
7. I understood the key drivers behind the predictions of my risk-profile
8. The prediction reasoning on the dashboard was clear about how the prediction of my risk-profile was made
9. The prediction reasoning was a useful insight
10. In future, I would like to see suggestions about what I can modify in my learning patterns which might result in better learning outcomes.

Responses to these questions were formulated on a five-point Likert scale (1: strongly disagree to 5: strongly agree). The total score is 100 and the maximum point for each question is 10. For each odd-numbered question, 1 was subtracted from their score while for each even-numbered question, 5 was subtracted from their value. These were added up; finally, this total score was multiplied by 2.5 to ensure that the maximum is 10 for each of the questions (Tulis & Stetson, 2004; Brooke, 2013).

6. Study Analysis and Findings

This section presents results from the statistical and survey analyses that measure the effects that SensEnablr had on students. The analyses seek to investigate both the impact of its usage on participants' behaviors after accessing the dashboard and its perceived effectiveness based on reported survey responses.

6.1 Dashboard analysis

The purpose of this analysis was to determine whether any significant changes in student engagement levels with the LMS had occurred in a period immediately following interactions with the dashboard (which was defined as 15 days). The paired sample t-test was conducted on data covering three 15-day before-and-after sequences covering the full range of the teaching semester in order to compare the online behavior effects of the participants after interacting with the dashboard.

The null hypothesis (H₀): $\mu_d = 0$, states that the difference in means between sample 1 (before using the dashboard) and sample 2 (after using the dashboard) is equal to 0. The alternate hypothesis (H_A): $\mu_d \neq 0$, states otherwise. The mean frequency of the participants engaging with Moodle was on average 115 times in the lead-up to the use of SensEnablr. This increased to an average of 213 times following the SensEnablr interactions. The t-test shows produced a statistical value of -2.14 with a p-value of .033. The p-value thus confirms that this constituted a significant difference of means at a .05 significance level allowing us to conclude that the online engagement with Moodle was different before and after interacting with the dashboard. However, this effect was not long-lasting. When the t-test was conducted to determine if the increased engagement with the LMS continued for students for a further 15-day period after the initial interaction with SensEnablr, the results confirmed that the differences in the means were not of a significant level.

6.2 Survey Analysis

The participants were asked to respond to a survey questionnaire for evaluating the dashboard's usability and its perceived effectiveness when they were halfway through the course. The purpose of this evaluation was to investigate whether students consider the dashboard useful in guiding them in making adjustments to their learning behaviors and if it motivated them to improve their learning outcomes. The outcome of the survey results is shown in Figure 6. The overall evaluation of the SUS questionnaire resulted in an average score of 70.5 points. In general, a score above 68 is considered to be positive (Brooke, 1996); therefore, it can be interpreted that from the standpoint of the students, the dashboard as a whole was perceived as helpful.

The survey responses corroborate the previous statistical results where a significant increase in engagement with the LMS was noted in response to interactions with the dashboard. The strongest and most positive response was registered on the question of whether or not SensEnablr motivated the students to continue learning. 67% responded as strongly agreeing with this assertion while a further 13% agreed with it. Therefore, an increase in engagement levels with the learning materials offered through the LMS can be assumed as a plausible effect arising from reported increases in motivation levels by the students' post-interaction with the LAD.

With respect to the predictive analytics features of the dashboard, 87% of the respondents either agreed or strongly agreed that displaying their risk profile positively affected their attitude toward their studies, with the remainder disagreeing. However, while the same percentage of respondents thought positively of the explanatory aspects of the model's reasoning, some further work in refining the dashboard is needed since two-thirds of the respondents did not agree that the dashboard was sufficiently

clear enough about what the exact drivers behind the predictions of their stated risk-profile were. Concerning the presentation of the prescriptive analytics features, 80% of the respondents were positive and indicated that they would like to see suggestions on which learning behaviors they should modify to improve learning outcomes. In considering the general issue of whether the SensEnablr satisfied the requirement of being both comprehensive and informative without inducing cognitive overload, two-thirds of the respondents responded as either agreeing or strongly agreeing that the users can easily understand the information at a glance.




Figure 6. Survey results

7. Findings

Despite the increased use of LADs by higher education institutions, this study finds that existing studies have not fully leveraged and integrated all available analytics capabilities. As such, no dashboard to date has brought together descriptive, predictive and prescriptive components; therefore, highlighting a gap that exists in current studies with regard to maximally extracting value from analytics technologies. Beyond pure academic interest, this research grounds the novel machine learning features implemented in SensEnablr within the existing theoretical frameworks and justifies their inclusion and value for supporting and enhancing learners and their outcomes. The descriptive analytics components, while unsophisticated, motivate learners to evaluate themselves by comparing their LMS levels of engagement with that of their peers thus drawing on the Social Comparison Theory. Meanwhile, the embedded ability for the learners to monitor their progress over time both with LMS engagement and assignment scores enables a greater degree of self-regulated learning to take place. The constructivist principles also come into play within SensEnablr, since learners are able to make sense of their performance through the assessment scores visualizations, from which they can identify gaps in their understanding, and be motivated to adjust their learning strategies accordingly. The decision to display to the students where their assessment score performance is situated within the context of their peers and what their next assessment score is likely to be (all things being equal), the learners are empowered to set better and more realistic goals which support the development of greater self-efficacy as outlined by the Social Cognitive Theory. Meanwhile, the predictive components which prognosticate a learner’s possible course outcome have the potential to initiate within learners a profound shift in their perspectives and beliefs through a confronting and critical self-reflection that challenges their existing assumptions about their learning capabilities as outlined by Transformative Learning. But, in assisting the learners to subsequently formulate new strategies as needed, the prescriptive features in SensEnablr then also enable the learners to develop new strategies with the additional insights that are offered.

Finally, while there is theoretical support for developing and using LADs in real settings, there is limited research investigating the overall effects of LAD usage on students. It is the purpose of this study to contribute towards addressing these gaps by investigating the effects it has had on students. Thus, in response to the first research question, this study explored

SensEnablr's effects on students through the prism of changes to the engagement levels with the LMS. Our statistical analyses confirmed that there was a significant improvement in the participants' level of engagement with Moodle before and after using the student dashboard. The effect was observed over the short-term engagement where the levels increased during the first 15 days following the usage of SensEnablr.

The second research question measured students' perceptions of the dashboard's effect on their learning. Generally, SensEnablr scored positively on its overall perception of effectiveness while scoring well on the usability component. Importantly, the survey results showed that two-thirds of the respondents felt that the insights about the learning activities from the dashboard motivated them to continue learning, which aligned with the statistical analysis showing that the students engaged with the LMS more after interacting with the dashboard. The survey responses showed that at least from this sample, the students were highly in favor of being shown predictive outputs of their risk profiles and were equally desirous of being offered clear prescriptive feedback through the LADs. However, further work needs to be conducted on making the transparency and understandability of the models and their outputs more comprehensible.

Overall, the integration of multiple frames of reference as defined by Jivet et al. (2017), and the attempt to utilize more fully all types of available analytics capabilities, have demonstrated promise and grounds for further refinement of future LADs. By anchoring SensEnablr within the reviewed theoretical frameworks, this work provides a comprehensive approach that ensures students' learning experiences are well-supported by established educational theories. Within the context of this experiment, it may be said that empowering the students with their own student-centered LAD plays a role in enhancing their awareness of their performance in comparison with that of their classmates; it can turn them into more responsible and competitive students who want to positively strive with their peers. Consequently, a student-centered analytical dashboard as presented in this study represents a possible blueprint and a step towards enhancing the state-of-the-art in learning analytics research.

In future work, the plan is to conduct a much larger rollout of the dashboard and to conduct analyses involving a significantly larger participant cohort. A particular emphasis in the subsequent work will be on quantifying the effects of the dashboard on learning outcomes in the form of course completion rates and grades, with a focus on the effects on at-risk students in particular. The aim is to also re-design aspects of the dashboard in order to make some explanatory components surrounding predictive and prescriptive capabilities more understandable admitting that there is room for improvement.

Furthermore, as a student observed their progress and engagement levels relative to other students, this may bring about psychological distress in the form of anxiety or discouragement if their expectations are not met. The dashboard may also bring some distress to high-risk students, who are considered to be on the path to non-completion or failure to meet learning outcomes. These questions need further research. The intention is to make a dashboard extension for instructors. Instructors could be provided detailed visuals on student performance to assist them in evaluating the online learning behavior of their students. By monitoring the online learning of the students through a dashboard, the instructors can support their students to perform better, especially those students considered at some risk of not accomplishing their learning objectives.

8. Study Limitation

Several limitations exist within this research. The sample size in this study is relatively low despite best efforts to recruit higher numbers. This naturally affects the confidence with which definitive conclusions can be made. Due to this, the differences in impact between high-performing and low-performing students which have been shown to exist in prior works could not effectively be investigated. Since causality is difficult to establish, it is not possible to claim with certainty that the increase in engagement levels of students who interacted with SensEnablr is due to this factor alone. There could be other factors at play that are connected with the courses and assessments or required LMS activities that might occur at different points in time during a semester which could be confounding factors. However, a mitigating factor to this is the fact that the study included participants from a number of different courses, each with its unique rhythm and schedule and the fact that the opt-in students tended to commence their usage of SensEnablr at different points in time. It is therefore the case that these factors helped attenuate any potential bias arising from the specific rhythms of individual courses. In future studies though, ideally, the participants would be drawn from a single large class and the changes in engagement levels of opt-in students could then be compared with opt-out students. Additionally, a further limitation of this study due to the small sample size was that the opt-in students were not large enough to randomly split into a control and treatment group which would mitigate some aspects surrounding the selection bias.

9. Conclusion

In this paper, a student-facing, learning analytics dashboard (SensEnablr) was presented together with an evaluation of its effectiveness to impact on learning for a group of students. The dashboard was developed for the purpose of prompting self-

reflection within students so as to trigger positive behavioral adjustments when necessary. A rich set of novel analytics insights were used and grounded within several educational theoretical frameworks. With this alignment, the dashboard's design enabled the students to self-evaluate their progress across the external frame (comparison with peers) and internal (comparison with their earlier performance and their goals) frames of reference. The dashboard incorporated into its capabilities a mixture of descriptive analytics displays, through to more sophisticated predictive outputs, as well as prescriptive analytics which offered students automated data-driven suggestions.

An evaluation analysis of the effectiveness of the dashboard was performed to investigate if interactions with the dashboard were associated with increased engagement levels with the institutional LMS together with an accompanying student survey. The result showed that the students did indeed exhibit higher engagement levels with the institutional learning platform and course resources immediately following interactions with the dashboard. Additionally, the survey respondents overwhelmingly indicated that the dashboard increased their learning motivation levels and that the sophisticated analytics insights were beneficial to their learning. Future work aims to expand both the size of the evaluation group as well as the length of the investigation in order to conduct a longitudinal study that can shed more light on the effects of this tool on academic attainment performance, course completion, and qualification retention rates.

Statement on Ethics

Informed consent was obtained from all study participants. All collected data was treated confidentially. This research has been approved by the university ethics board (NOR 19/55).

Declaration of Conflicting Interest

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The authors declared no financial support for the research, authorship, and/or publication of this article.

References

- Afzaal, M., Nouri, J., Zia, A., Papapetrou, P., Fors, U., Wu, Y., Li, X., & Weegar, R. (2021). Explainable AI for data-driven feedback and intelligent action recommendations to support students self-regulation. *Frontiers in Artificial Intelligence*, 4. <https://doi.org/10.3389/frai.2021.723447>
- Ahadi, A., Lister, R., Haapala, H., & Vihavainen, A. (2015). Exploring machine learning methods to automatically identify students in need of assistance. In *Proceedings of the eleventh annual international conference on international computing education research* (pp. 121-130). Omaha, Nebraska, USA: ACM. <https://doi.org/10.1145/2787622.2787717>
- Ali, L., Asadi, M., Gašević, D., Jovanović, J., & Hatala, M. (2013). Factors influencing beliefs for adoption of a learning analytics tool: An empirical study. *Computers & Education*, 62, 130-148. <https://doi.org/10.1016/j.compedu.2012.10.023>
- Aljohani, N. R., Daud, A., Abbasi, R. A., Alowibdi, J. S., Basher, M., & Aslam, M. A. (2019). An integrated framework for course adapted student learning analytics dashboard. *Computers in Human Behavior*, 92, 679-690. <https://doi.org/10.1016/j.chb.2018.03.035>
- Arnold, K. E., & Pistilli, M. D. (2012, April). Course signals at Purdue: Using learning analytics to increase student success. In *Proceedings of the 2nd international conference on learning analytics and knowledge* (pp. 267-270). <https://doi.org/10.1145/2330601.2330666>
- Bandura, A. (1986). The explanatory and predictive scope of self-efficacy theory. *Journal of social and clinical psychology*, 4(3), 359-373. <https://doi.org/10.1521/jscp.1986.4.3.359>
- Baneres, D., Rodriguez, M. E., & Serra, M. (2019). An early feedback prediction system for students at-risk within a first-year higher education course. *IEEE Transactions on Learning Technologies*. <https://doi.org/10.1109/TLT.2019.2912167>
- Bañeres, D., Rodríguez, M. E., Guerrero-Roldán, A. E., & Karadeniz, A. (2020). An early warning system to detect at-risk students in online higher education. *Applied Sciences*, 10(13), 1–28. <https://doi.org/10.3390/app10134427>
- Bodily, R., & Verbert, K. (2017). Review of research on student-facing learning analytics dashboards and educational recommender systems. *IEEE Transactions on Learning Technologies*, 10(4), 405–418. <https://doi.org/10.1109/TLT.2017.2740172>

- Bodily, R., Ikahihifo, T. K., Mackley, B., & Graham, C. R. (2018). The design, development, and implementation of student-facing learning analytics dashboards. *Journal of Computing in Higher Education*, 30(3), 572–598. <https://doi.org/10.1007/s12528-018-9186-0>
- Brooke, J. (1996). SUS: A "quick and dirty" usability scale. In P. W. Jordan, B. Thomas, B. A. Weerdmeester, & A. L. McClelland (Eds.), *Usability Evaluation in Industry*. London: Taylor and Francis
- Brooke, J. (2013). SUS: A Retrospective. *Journal of usability studies*, 29–40.
- Breiman, L. Random forests. *Mach. Learn.* 2001, 45, 5–32. Chatti, M. A., Muslim, A., Guliani, M., & Guesmi, M. (2020). The LAVA model: learning analytics meets visual analytics. https://doi.org/10.1007/978-3-030-47392-1_5
- Cranton, P. (2006). *Understanding and promoting transformative learning: A guide for educators of adults*.
- Domingos, P.; Pazzani, M. On the Optimality of the Simple Bayesian Classifier under Zero-One Loss. *Mach. Learn.* 1997, 29, 103–130.
- Dorogush, A. V., Ershov, V., & Gulin, A. (2018). CatBoost: gradient boosting with categorical features support. <http://arxiv.org/abs/1810.11363>
- Duan, X., Wang, C., & Rouamba, G. (2022). Designing a Learning Analytics Dashboard to Provide Students with Actionable Feedback and Evaluating Its Impacts. In *CSEdu (2)* (pp. 117-127).
- Fleur, D. S., van den Bos, W., & Bredeweg, B. (2020). Learning analytics dashboard for motivation and performance. https://doi.org/10.1007/978-3-030-49663-0_51
- Hajian, S., Bonchi, F., & Castillo, C. (2016, August). Algorithmic bias: From discrimination discovery to fairness-aware data mining. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 2125-2126). <https://doi.org/10.1145/2939672.2945386>
- Han, J., Kim, K. H., Rhee, W., & Cho, Y. H. (2021). Learning analytics dashboards for adaptive support in face-to-face collaborative argumentation. *Computers & Education*. <https://doi.org/10.1016/j.compedu.2020.104041>
- Hechenbichler, K.; Schliep, K. Weighted k-Nearest-Neighbor Techniques and Ordinal Classification Projektpartner Weighted k-Nearest-Neighbor Techniques and Ordinal Classification. 2004. Available online: <http://epub.ub.uni-muenchen.de/> (accessed on 4 October 2021)
- Hellings, J., & Haelermans, C. (2020). The effect of providing learning analytics on student behaviour and performance in programming: A randomised controlled experiment. *Higher Education*. <https://doi.org/10.1007/s10734-020-00560-z>
- Hosmer, D.W.; Lemeshow, S. *Applied Logistic Regression*; John Wiley & Sons: New York, NY, USA, 2000.
- Howell, J. A., Roberts, L. D., & Mancini, V. O. (2018). Learning analytics messages: Impact of grade, sender, comparative information and message style on student affect and academic resilience. *Computers in Human Behavior*, 89, 8-15. <https://doi.org/10.1016/j.chb.2018.07.021>
- Islam, M.R.; Ahmed, M.U.; Barua, S.; Begum, S. A systematic review of explainable artificial intelligence in terms of different application domains and tasks. *Appl. Sci.* 2022, 12, 1353. <https://doi.org/10.3390/app12031353>
- Jayaprakash, S. M., Moody, E. W., Lauría, E. J. M., Regan, J. R., & Baron, J. D. (2014). Early alert of academically at-risk students: An open source analytics initiative. *Journal of Learning Analytics*, 1, 1–42. <https://doi.org/10.18608/jla.2014.11.3>
- Jivet, I., Scheffel, M., Drachler, H., & Specht, M. (2017, September). Awareness is not enough: Pitfalls of learning analytics dashboards in the educational practice. In *European conference on technology enhanced learning* (pp. 82-96). Springer, Cham. https://doi.org/10.1007/978-3-319-66610-5_7
- Jivet, I., Scheffel, M., Specht, M., & Drachler, H. (2018, March). License to evaluate: Preparing learning analytics dashboards for educational practice. In *Proceedings of the 8th international conference on learning analytics and knowledge* (pp. 31-40). <https://doi.org/10.1145/3170358.3170421>
- Khosravi, H., Shabaninejad, S., Bakharia, A., Sadiq, S., Indulska, M., & Gasevic, D. (2021). Intelligent Learning Analytics Dashboards: Automated Drill-Down Recommendations to Support Teacher Data Exploration. *Journal of Learning Analytics*, 8(3), 133-154. <https://doi.org/10.18608/jla.2021.7279>
- Kim, J., Jo, I. H., & Park, Y. (2016). Effects of learning analytics dashboard: analyzing the relations among dashboard utilization, satisfaction, and learning achievement. *Asia Pacific Education Review*, 17(1), 13-24. <https://doi.org/10.1007/s12564-015-9403-8>
- Kitchenham, A. (2008). The evolution of John Mezirow's transformative learning theory. *Journal of transformative education*, 6(2), 104-123 <https://doi.org/10.1177/1541344608322678>
- Knight, D., Knight, B., Pearson, M., & Quintana, M. (2018). *Microsoft Power BI Quick Start Guide: Build dashboards and visualizations to make your data come to life*. Packt Publishing Ltd.
- Kokoç, M., & Altun, A. (2021). Effects of learner interaction with learning dashboards on academic performance in an e-

- learning environment. *Behaviour & Information Technology*. <https://doi.org/10.1080/0144929X.2019.1680731>
- Lim, L., Dawson, S., Joksimovic, S., & Gašević, D. (2019, March). Exploring students' sensemaking of learning analytics dashboards: Does frame of reference make a difference?. In *Proceedings of the 9th international conference on learning analytics & knowledge* (pp. 250-259) <https://doi.org/10.1145/3303772.3303804>
- Lu, C., & Cutumisu, M. (2022). Online engagement and performance on formative assessments mediate the relationship between attendance and course performance. *International journal of educational technology in higher education*, 19(1), 1-23. <https://doi.org/10.1186/s41687-018-0043-8>
- Lundberg, S.M.; Allen, P.G.; Lee, S.-I. A Unified Approach to Interpreting Model Predictions. Available online: <https://github.com/slundberg/shap> (accessed on 12 October 2021).
- Lepenioti, K., Bousdekis, A., Apostolou, D., & Mentzas, G. (2020). Prescriptive analytics: Literature review and research challenges. *International Journal of Information Management*, 50, 57-70. <https://doi.org/10.1016/j.ijinfomgt.2019.04.003>
- Marsh, H. W. (1986). Verbal and math self-concepts: An internal/external frame of reference model. *American educational research journal*, 23(1), 129-149 <https://doi.org/10.3102/00028312023001129>
- Marsh, H. W., Trautwein, U., Lüdtke, O., Köller, O., & Baumert, J. (2005). Academic self-concept, interest, grades, and standardized test scores: Reciprocal effects models of causal ordering. *Child development*, 76(2), 397-416. <https://doi.org/10.1111/j.1467-8624.2005.00853.x>
- Mathrani, A., Susnjak, T., Ramaswami, G., & Barczak, A. (2021). Perspectives on the challenges of generalizability, transparency and ethics in predictive learning analytics. *Computers and Education Open*, 2, 100060. <https://doi.org/10.1016/j.caeo.2021.100060>
- Mezirow, J. (1997). Transformative learning: Theory to practice. *New directions for adult and continuing education*, 1997(74), 5-12. <https://doi.org/10.1177/1541344608322678>
- Namoun, A., & Alshantqiti, A. (2020). Predicting student performance using data mining and learning analytics techniques: A systematic literature review. *Applied Sciences*, 11(1), 237,1-28. <https://doi.org/10.3390/app11010237>
- Piaget, J. (1952). *The origins of intelligence in children* (M. Cook, Trans.). International Universities Press. <https://doi.org/10.1037/11494-000>
- Pintrich, P. R. (2000). The role of goal orientation in self-regulated learning. In *Handbook of self-regulation* (pp. 451-502). Academic Press. <https://doi.org/10.1016/B978-012109890-2/50043-3>
- de Quincey, E., Briggs, C., Kyriacou, T., & Waller, R. (2019). Student centred design of a learning analytics system. In: *Proceedings of the 9th international conference on learning analytics & knowledge* (pp 353–362). <https://doi.org/10.1145/3303772.3303793>
- Rets, I., Herodotou, C., Bayer, V., Hlosta, M., & Rienties, B. (2021). Exploring critical factors of the perceived usefulness of a learning analytics dashboard for distance university students. *International Journal of Educational Technology in Higher Education*, 18(1), 1-23. <https://doi.org/10.1186/s41239-021-00284-9>
- Ribeiro, M.T.; Singh, S.; Guestrin, C. Why Should I Trust You? In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Francisco, CA, USA, 13–17 August 2016; Association for Computing Machinery: New York, NY, USA; pp. 1135–1144. <https://doi.org/10.1145/2939672.2939778>
- Ribeiro, M.T.; Singh, S.; Guestrin, C. Anchors: High-Precision Model-Agnostic Explanations. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, New Orleans, LA, USA, 2–7 February 2018; pp. 1527–1535. <https://doi.org/10.1609/aaai.v32i1.11491>
- Rienties, B., Herodotou, C., Olney, T., Schencks, M., & Boroowa, A. (2018). Making sense of learning analytics dashboards: A technology acceptance perspective of 95 teachers. *International Review of Research in Open and Distributed Learning*, 19(5). <https://doi.org/10.19173/irrodl.v19i5.3493>
- Roberts, L. D., Howell, J. A., & Seaman, K. (2017). Give me a customizable dashboard: Personalized learning analytics dashboards in higher education. *Technology, Knowledge and Learning*, 22(3), 317-333. <https://doi.org/10.1007/s10758-017-9316-1>
- Ryan, T., & Henderson, M. (2018). Feeling feedback: students' emotional responses to educator feedback. *Assessment & Evaluation in Higher Education*, 43(6), 880-892. <https://doi.org/10.1080/02602938.2017.1416456>
- Schwendimann, B. A., Rodriguez-Triana, M.J., Vozniuk, A., Prieto, L.P., Boroujeni, M.S., Holzer, A., Gillet, D. & Dillenbourg, P. (2016). Perceiving learning at a glance: A systematic literature review of learning dashboard research. *IEEE Transactions on Learning Technologies*, 10(1), 30-41. <https://doi.org/10.1109/TLT.2016.2599522>
- Shahiri, A. M., & Husain, W. (2015). A review on predicting student's performance using data mining techniques. *Procedia Computer Science*, 72, 414-422. <https://doi.org/10.1016/j.procs.2015.12.157>
- Skaalvik, E. M., & Skaalvik, S. (2004). Self-concept and self-efficacy: A test of the internal/external frame of reference

- model and predictions of subsequent motivation and achievement. *Psychological reports*, 95(3_suppl), 1187-1202. <https://doi.org/10.2466/pr0.95.3f.1187-1202>
- Stahl, G., Koschmann, T., & Suthers, D. (2006). Computer-supported collaborative learning: An historical perspective. In Sawyer, R. K. (Ed.), *Cambridge handbook of the learning sciences* (pp. 409–426). Cambridge: Cambridge University Press
- Susnjak, T., Ramaswami, G. S., & Mathrani, A. (2022). Learning analytics dashboard: a tool for providing actionable insights to students. *International Journal of Educational Technology in Higher Education*, 19(1), 1-23. <https://doi.org/10.1186/s41239-021-00313-7>
- Susnjak, T. (2023). Beyond Predictive Learning Analytics Modelling and onto Explainable Artificial Intelligence with Prescriptive Analytics and ChatGPT. *International Journal of Artificial Intelligence in Education*, 1-31 <https://doi.org/10.1007/s40593-023-00336-3>
- Teasley, S. D. (2017). Student facing dashboards: One size fits all?. *Technology, Knowledge and Learning*, 22(3), 377-384. <https://doi.org/10.1007/s10758-017-9314-3>
- Tullis, T. S., & Stetson, J. N. (2004). A Comparison of Questionnaires for Assessing Website Usability. *Proceedings of UPA 2004 Conference*, (pp. 1-12). Minneapolis, Minnesota
- Valle, N.; Antonenko, P.; Valle, D.; Sommer, M.; Huggins-Manley, A.C.; Dawson, K.; Kim, D.; Baiser, B. Predict or describe? How learning analytics dashboard design influences motivation and statistics anxiety in an online statistics course. *Educ. Technol. Res. Dev.* 2021, 69, 1405–1431. <https://doi.org/10.1007/s11423-021-09998-z>
- Verbert, K., Duval, E., Klerkx, J., Govaerts, S., & Santos, J. L. (2013). Learning analytics dashboard applications. *American Behavioral Scientist*, 57(10), 1500–1509. <https://doi.org/10.1177/0002764213479363>
- Verbert, K., Ochoa, X., De Croon, R., Dourado, R. A., & De Laet, T. (2020, March). Learning analytics dashboards: The past, the present and the future. In *Proceedings of the tenth international conference on learning analytics & knowledge* (pp. 35-40) <https://doi.org/10.1145/3375462.3375504>.
- Vygotsky, L. S., & Cole, M. (1978). *Mind in society: Development of higher psychological processes*. Harvard university press.
- Winne, P. H., & Hadwin, A. (1998). Studying as Self-Regulated Learning. In D. J. Hacker, & J. Dunlosky (Eds.), *Metacognition in Educational Theory and Practice* (pp. 277-304), Routledge.
- Wise, A. F. (2014, March). Designing pedagogical interventions to support student use of learning analytics. In *Proceedings of the fourth international conference on learning analytics and knowledge* (pp. 203-211). <https://doi.org/10.1145/2567574.2567588>
- Y. Li, J. Gou, and Z. Fan, “Educational data mining for students’ performance based on fuzzy C-means clustering,” *The Journal of Engineering*, vol. 2019, no. 11, pp. 8245–8250, Nov. 2019, <https://doi.org/10.1049/joe.2019.0938>
- Yu, T., & Jo, I. H. (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*. Indianapolis, IN, USA. <https://doi.org/10.1145/2567574.2567594>
- Zacharis, N. Z. (2015). A multivariate approach to predicting student outcomes in web-enabled blended learning courses. *The Internet and Higher Education*, 27, 44-53. <https://doi.org/10.1016/j.iheduc.2015.05.002>
- Zimmerman, D. W. (1997). Teacher’s corner: A note on interpretation of the paired-samples t test. *Journal of Educational and Behavioral Statistics*, 22(3), 349-360 <https://doi.org/10.3102/1076998602200334>
- Zimmerman, B. J. (2000). Attaining self-regulation: A social cognitive perspective. In *Handbook of self-regulation* (pp. 13-39). Academic press. <https://doi.org/10.1016/B978-012109890-2/50031-7>.
- Zimmerman, B. J., & Schunk, D. H. (Eds.). (2001). *Self-regulated learning and academic achievement: Theoretical perspectives*. Routledge, <https://doi.org/10.1007/978-1-4612-3618-4>.

Chapter 9

Conclusion, Limitations and Future

Research Directions

9.1 Conclusions

To date, the major focus in LAD has been on descriptive, and predictive analytics. Existing studies in the educational domain mainly target the delivery of model's predicted results to students and do not exhibit which aspects of learners' behaviors are driving these predictions towards prognosticated outcomes. Also, they do not provide any recommendations to students on what learning adjustments could result in more positive predictions for them based on their observed behavioral learning patterns. Taking these gaps into account, this research study has demonstrated a LAD design and its implementation, that not only leverages descriptive analytics, but which further integrates different types of machine learning algorithms in a way that enables both predictive and prescriptive analytics.

After a thorough observation of background information on the current state-of-the art and gaps in LADs six research questions have been posed. The main objectives were set based on the proposed research questions to develop a sophisticated LAD which integrates all forms of analytics: descriptive, predictive, and prescriptive.

This study was initiated by conducting a systematic literature review to identify gaps and challenges in the current research into the utilization of predictive analytics tools among existing LADs. The systematic literature review provided an in-depth view of the types of data being used, which machine learning/data mining methods have been used for making predictions and how these were being presented to different stakeholders (student advisors, instructors and students). Moreover, the literature review revealed that the predictive analytics tools among existing LADs were not being leveraged fully. Prior studies that have utilized predictive analytics only go as far as identifying the at-risk students and have not employed model interpretation or explainability capabilities. This limits the ability of LADs to offer data-driven prescriptive advice to students that can guide them on appropriate learning adjustments for more favorable outcomes. This limitation in integrating all three analytics (i.e., descriptive,

predictive, and prescriptive) to provide a more meaningful representation of user interactions on LADs provided the direction for moving this study forward.

The first research question that was posed was to determine how a generic, course-agnostic predictive model can be developed to predict students' outcomes across disparate courses. Accordingly, many experiments were conducted to generate a generic and course-agnostic predictive model for identifying the at-risk students across a wide variety of courses. The purpose was to overcome the limitations of building multiple models, where each model is tightly coupled with the specific attributes of different courses. Students' data from LMS (Moodle) log along with demographic information and assignment scores from various semesters were used to predict the outcome after every week into two classes, namely at-risk and not-at-risk of failing the course. These experiments were carried out using the commonly used Random Forest, Naïve Bayes, Logistic Regression, and k-Nearest Neighbors algorithms, as well as the recently developed CatBoost algorithm. Across several performance metrics gathered from the experiments, a course-agnostic model built using CatBoost was considered most viable for the identifying the at-risk students. This model provided the best results and had the potential to reduce academic failure rates through early predictions and interventions. However, after developing any predictive model, it is important to provide data-driven suggestions to students who are at risk of non-completions or other sub-optimal outcomes. This will help to build more trust and accountability of the model among the users. Our review revealed that while prior LAD studies have predicted students' performance, they have not explained the reasoning behind the predictive decisions. This brings the focus to the second and third research questions, that is, how can explainable ML approaches be leveraged to provide effective and human-understandable reasoning behind their conclusions about a student's academic performance, and further, how can automatic data-driven feedback and actionable recommendations be derived? Henceforth, ML methods based on real-world data

that not only identified trends and predicted learning outcomes, but they also provided feedback of customized actionable suggestions to students and other stakeholders. By providing students with feedback not only alerted them to a potential problem but also explained to them why the alert was raised and what remedial steps can be taken by the user to avert the problem. Many experimentations were conducted using anchors to identify the influencing factors in the predictive models and then to explain to students how the model works and how the predictions have been generated in an easily understandable way. This in turn helps students understand their role and make good learning behavioral decisions. In addition, counterfactuals too were generated to provide us with alternate suggestions for overcoming unfavorable predictions. Next, the counterfactual explanations have been induced and translated into a human-readable format to provide students with clear, precise, and actionable suggestions on what adjustments to their learning behavior and performance are most likely to assist them in realizing positive outcomes.

After developing various analytics technologies, the next step involved building a dashboard plug-in by integrating the technologies and deploying them as a dashboard extension for some ongoing courses. This helped in answering the fourth research question (Determine how the identified model can be projected to students so that the students can view their current academic performance as putting them at low risk, or high risk of not meeting their learning outcomes). This research question was posed as a consequence of gaps identified in existing dashboards from the systematic literature review. That is, the review revealed that most of the existing dashboards merely employed surface-level descriptive analytics, while only a few had gone beyond and used predictive analytics. In response to the identified gaps in the prior literature, this dashboard design was proposed and is the first of its kind in terms of breadth of analytics it integrates comprising all three analytics, namely, descriptive, predictive, and prescriptive analytics.

Finally, participants were recruited from various disciplines, and they were informed about the nature and purpose of this study. Only those students who opted-in by signing the participant consent form were considered as participants. This dashboard access was given to all participants four weeks into the start of the semester and access was enabled until the end of the semester. The last two research questions (i.e., Does the student-facing analytical dashboard result in more student engagement with the institutional LMS? and What are the students' perceptions of the dashboard's effect on their learning performance?) were aimed to gauge the effectiveness of the proposed LAD among student users. Participants were surveyed to gain insight on their perspectives, such as whether the LAD supported them in their learning journey over the semester. Using statistical analysis, the effectiveness of the LAD and the association between its use and level of course performance were studied. A mixed-methods approach that looked at online student behavior after accessing the dashboard and survey questionnaire that gauged students' perceptions of the dashboard was undertaken. Our results showed that a change in online behavior of students after accessing the dashboard and overall, the usability was also perceived as good.

9.2 Limitations and Future Research Directions

This section briefly lists the study limitations and future research directions in the field of learning analytics dashboards.

First, this study was conducted during the pandemic period, which was a period of some uncertainty as institutions were forced to adopt online course delivery modes. Thus, the results obtained from this study are influenced to some extent by the pandemic restrictions and therefore may not be very precise. Secondly, the sample size was relatively small, since the sample comprised solely of student participants who had opted in voluntarily. Moreover, the impact of the dashboard usage could not be conducted by dividing the participants into

treatment and control groups, where the students in the treatment group would have been given access to the dashboard. This would have helped attribute any differences observed in the results between the two groups more precisely and would have given more insight on the effectiveness of the dashboard. Many existing studies [29-31] have tested the effectiveness of their dashboards on student final outcomes, to ascertain whether the dashboard helped the students in achieving their desired result or not. Unfortunately, in this study, the final performance score of the participants were not made available to us, hence, the impact of the dashboard on the students' outcomes could not be studied. Though the accuracy obtained for the predictive model was more than 73%, this could have been improved further by adding more features, or by enhancing the training dataset with use of resampling methods for imbalanced classes.

To make the LAD more useful for the students, and to perform better in a course, additional features such as library visits to understand more learning patterns, such as how often a student is accessing the library resources or is using the institutional WiFi for attending courses could be used. These would center around the insights of effective learning and guide students in their development as life-long learners

It is important to understand the impact of the LAD on different student populations. Literature has highlighted some gaps regarding student responses to dashboards. While some studies show dashboards can have an overall positive impact, there is not much information about which students are more positively impacted or whether there are students that are negatively motivated by dashboards, or why are students influenced in any manner. As the student sees their progress and engagement levels relative to other students, this may bring about psychological harm in the form of anxiety or discouragement if the student's expectations are not met. Or, the dashboard may bring some distress to a high-risk student, who is considered

to be on the path to non-completion or failure to meet learning outcomes. These questions remain unanswered and need future research.

Another research direction would be to make a dashboard extension for teachers. Teachers could be provided detailed visuals on student performance to assist them in evaluating online learning behavior of their students. By monitoring the online learning of the students through a dashboard, the teachers can support their students to perform better, especially those students considered at some risk of not accomplishing their learning objectives.

On a concluding note, this thesis affirms the utility of having a robust analytic environment that supports all three kinds of analytics (descriptive, predictive, and prescriptive). Learning Analytics Dashboards are being seen as the next learning and teaching tool to enhance student learning by supporting them in understanding their learning behaviors. This thesis is a step forward in showcasing the design, implementation and evaluation of a student-facing dashboard that was set up in a real-world setting.

References

- [1] L. P. Macfadyen and S. Dawson, "Mining LMS data to develop an "early warning system" for educators: A proof of concept.," *Computers & education*, pp. 588-599, 2010. <https://doi.org/10.1016/j.compedu.2009.09.008>
- [2] N. Cavus, "Distance Learning and Learning Management Systems," *Procedia-Social and Behavioral Sciences*, pp. 872-877, 2015. <https://doi.org/10.1016/j.sbspro.2015.04.611>
- [3] C. Romero and S. Ventura, "Educational Data Mining: A Review of the State of the Art," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, pp. 601-618, 2010. DOI: 10.1109/TSMCC.2010.2053532
- [4] Siemens, "Learning analytics: The emergence of a discipline," *American Behavioral Scientist*, pp. 1380-1400, 2013. <https://doi.org/10.1177/000276421349885>
- [5] T. E. Lias and T. Elias, "Learning Analytics: The Definitions, the Processes, and the Potential," pp. 1-23, 2011.
- [6] L. Ali, M. Asadi, D. Gašević, J. Jovanović and M. Hatala, "Factors influencing beliefs for adoption of a learning analytics tool: An empirical study," *Computers & Education*, pp. 130-148, 2013. <https://doi.org/10.1016/j.compedu.2012.10.023>
- [7] G. Siemens and D. Gasevic, "Guest editorial-learning and knowledge," *Journal of Educational Technology & Society*, pp. 1-2, 2012.
- [8] S. Dawson, D. Gašević, G. Siemens and S. Joksimovic, "Current state and future trends: A citation network analysis of the learning analytics field," in *In Proceedings of the fourth international conference on learning analytics and knowledge*, Indianapolis, 2014. <https://doi.org/10.1145/2567574.2567585>
- [9] D. Gašević, S. Dawson and G. Siemens, "Let's not forget: Learning analytics are about learning" *TechTrends*, pp. 64-71, 2015.
- [10] N. Sclater, A. Peasgood and J. Mullan, " Learning analytics in higher education.," Jisc, 2016. DOI 10.1007/978-3-319-52977-6_1
- [11] M. Hanna, "Data mining in the e-learning domain," *Campus-Wide Information Systems*, pp. 29-34, 2004. <https://doi.org/10.1108/10650740410512301>
- [12] W. He, "Examining students' online interaction in a live video streaming environment using data mining and text mining," *Computers in Human Behavior*, pp. 90-102, 2013. <https://doi.org/10.1016/j.chb.2012.07.020>
- [13] G. Lust, J. Elen and G. Clarebout, "Students' tool-use within a web enhanced course: Explanatory mechanisms of students' tool-use pattern," *Computers in Human Behavior*, pp. 2013-2021, 2013. <https://doi.org/10.1016/j.chb.2013.03.014>
- [14] J. López-Zambrano, J. A. Lara and C. Romero, "Towards Portability of Models for Predicting Students' Final Performance in University Courses Starting from Moodle Logs.," *Applied Sciences*, pp. 1-23, 2020. <https://doi.org/10.3390/app10010354>
- [15] G. Ramaswami, T. Susnjak and A. Mathrani, "Supporting Students' Academic Performance Using Explainable Machine Learning with Automated Prescriptive Analytics," *Big Data and Cognitive Computing*, pp. 1-14, 2022. <https://doi.org/10.3390/bdcc6040105>

- [16] C. Vieira, P. Parsons and V. Byrd, "Visual learning analytics of educational data: A systematic literature review and research agenda," *Computers & Education*, pp. 119-135, 2018.
<https://doi.org/10.1016/j.compedu.2018.03.018>
- [17] R. Bodily and K. Verbert, "Review of research on student-facing learning analytics dashboards and educational recommender systems," *IEEE Transactions on Learning Technologies*, pp. 405-418, 2017.
DOI: 10.1109/TLT.2017.2740172
- [18] Y. Park and I. H. Jo, "Development of the learning analytics dashboard to support students' learning performance," *Journal of Universal Computer Science*, pp. 110-133, 2015.
- [19] G. Sedrakyan, J. Malmberg, K. Verbert, S. Järvelä and P. A. Kirschner, "Linking learning behavior analytics and learning science concepts: Designing a learning analytics dashboard for feedback to support learning regulation.," *Computers in Human Behavior*, pp. 1-15, 2020.
<https://doi.org/10.1016/j.chb.2018.05.004>
- [20] S. Few, "Information dashboard design: The effective visual communication of data," O'Reilly Media, Inc., 2006.
- [21] B. A. Schwendimann, M. J. Rodríguez-Triana, A. Vozniuk, L. P. Prieto, M. S. Boroujeni, A. Holzer and P. Dillenbourg, "Perceiving learning at a glance: A systematic literature review," *IEEE Transactions on Learning Technologies*, pp. 30-41, 2016. DOI: 10.1109/TLT.2016.2599522
- [22] B. A. Schwendimann, M. J. Rodríguez-Triana, A. Vozniuk, L. P. Prieto, M. S. Boroujeni, A. Holzer and P. Dillenbourg, "Understanding learning at a glance: An overview of learning," in *In Proceedings of the sixth international conference on learning analytics & knowledge*, Edinburgh, 2016.
<http://dx.doi.org/10.1145/2883851.2883930>
- [23] A. L. Dyckhoff, D. Zielke, M. Bültmann, M. A. Chatti and U. Schroeder, "Design and Implementation of a Learning Analytics Toolkit for Teachers," *Educational Technology & Society*, pp. 58-76, 2012.
- [24] S. Lange, H. Schumann, W. Müller and D. Krömker, "Problem-oriented visualisation of multi-dimensional data sets," in *In Proceedings of the International Symposium and Scientific Visualization*, 1995.
- [25] B. Shneiderman, "The Eyes Have It: A Task by Data Type Taxonomy of Information Visualization," in *IEEE Symposium on Visual Languages*, 1996.
- [26] S. Stoltzman, "Data Visualization – Part 3," 11 February 2018. [Online]. Available:
https://opendatascience.com/data-visualization-part-3/?utm_content=67206365&utm_medium=social&utm_source=googleplus.
- [27] G. Sedrakyan, E. Mannens and K. Verbert, "Linking dashboard design and data visualization concepts" *Journal of Computer Languages*, pp. 19-38, 2019.
- [28] J. L. Santos, S. Govaerts, K. Verbert and E. Duval, "Goal-oriented visualizations of activity tracking: a case study with engineering students" in *In Proceedings of the 2nd international conference on learning analytics and knowledge*, 2012. <https://doi.org/10.1145/2330601.2330639>
- [29] Bañeres, D., Rodríguez, M. E., Guerrero-Roldán, A. E., & Karadeniz, A. (2020). An early warning system to detect at-risk students in online higher education. *Applied Sciences*, 10(13), 1–28.
<https://doi.org/10.3390/app10134427>
- [30] Fleur, D. S., van den Bos, W., & Bredeweg, B. (2020). Learning analytics dashboard for motivation and performance. https://doi.org/10.1007/978-3-030-49663-0_51

[31] Kim, J., Jo, I. H., & Park, Y. (2016). Effects of learning analytics dashboard: analyzing the relations among dashboard utilization, satisfaction, and learning achievement. *Asia Pacific Education Review*, 17(1), 13-24. DOI 10.1007/s12564-015-9403-8

Appendices

Appendix 1

The contents of this appendix are from the following article. In accordance with the IEEE republishing policy, any full text that has been included, is the unmodified accepted article.

@2020 IEEE. Reprinted, with permission, from: Ramaswami, G. S., Susnjak, T., & Mathrani, A. (2019, December). Capitalizing on Learning Analytics Dashboard for Maximizing Student Outcomes. In 2019 IEEE Asia-Pacific Conference on Computer Science and Data Engineering (CSDE) (pp. 1-6). IEEE.,

DOI: 10.1109/CSDE48274.2019.9162357

The IEEE open access policy, permits the use, sharing, adaptation, distribution and reproduction in any medium or format as long as appropriate credit goes to the authors. IEEE does not endorse any of Massey University's products or services. Internal or personal use of this material is permitted. If interested in reprinting/republishing IEEE copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribution, please go to: http://www.ieee.org/publications_standards/publications/rights/rights_link.html to learn how to obtain a license from the correct link.

Capitalizing on Learning Analytics Dashboard for Maximizing Student Outcomes

Gomathy Suganya Ramaswami
*School of Natural and Computational
Studies*
Massey University
Auckland, New Zealand
ramaswamigomathy@gmail.com

Teo Susnjak
*School of Natural and Computational
Studies*
Massey University
Auckland, New Zealand
T.Susnjak@massey.ac.nz

Anuradha Mathrani
*School of Natural and Computational
Studies*
Massey University
Auckland, New Zealand
A.S.Mathrani@massey.ac.nz

Abstract— With ongoing advancements in technology enabled learning, an opportunity has risen for educators to enhance student learning with use of Learning Analytics. Educational institutions are using Learning Analytics Dashboards (LAD) to provide students with timely and personalized feedback in visual format. LAD use advanced analytical capabilities that capitalize on online learner activities that are extracted from log files. They provide data-driven insights on current learning contexts and inform management on any learning intervention strategies that may be needed to support students in achieving their learning outcomes. Besides students, the perspective of instructors too is considered. Using easy-to-read graphical reporting formats (e.g. line graphs, scatter plot, bar charts, etc.), the LAD reveals a consolidated view of how online learning is taking place. In this manner, a snapshot depicting details of student learning patterns can enable instructors to monitor their students' learning strategies. At the same time, the LAD assists students too by providing them with a personalized environment to help them engage better with the learning outcomes. Therefore, LAD is increasingly used as a pedagogical approach to motivate students and help build their self-learning capacity. In this study, we propose to develop a real-time dashboard that pulls online student data from various sources including a learning management system (Moodle), University's library and from the student management system (SMS) that is used by the staff.

Keywords—learning dashboard, feedback, student feedback system, visualization tool, dashboard

I. INTRODUCTION

Recent trends in learning analytics have been to use mined learning data for developing automated interventions that can be directly fed to reports and displayed in graphical formats [1]; thereby enabling students and educators to apply data driven insights in their decision making. This move has led to the development of dashboards for making visualizations on key performance indicators that can inform educators and learners on any issues that may have been identified by the data analysis techniques [2]. Visualization tools can summarize large quantities of data from basic charts to graphical indicators. Therefore, visualization plays an essential role in empowering users with relevant information about the overall learning process. The main aim is to provide different end-users (students, teachers, study advisers, administrators etc.) with a Learning Analytics dashboard (LAD) that highlights informative learning patterns [3]. Malik [15] notes that dashboards are user specific, implying that every user has specific requirements which a dashboard needs to satisfy.

An educational dashboard has various roles and values that have to be considered. To begin with, LAD enables

educators to know students' learning status in real time and in a scalable way. Student activity patterns reported in the dashboard are specifically purposed in online learning environments. First, the students and educators are physically isolated. Thus, a LAD can be used to provide a personalized learner centered view to students and an overview of student performances to educators. Second, students' self-knowledge can be improved by reviewing their personal learning status and history through the dashboard [4]. Most importantly, information displayed in the dashboard can lead to more intelligent decision making. For instance, dashboard can aid in identifying the at-risk students or in predicting the high performers, as well as communicating suggestions, giving personalized feedback and showcasing guidelines to support students.

An extensive range of dashboards have been designed, developed and evaluated that can track student activities in web-based learning environments [5]. These dashboards vary depending upon the type of data used for analysis, the visuals portrayals used for displaying information and the evaluation techniques used. While many data variables can be rendered on dashboards, the visualizations should be intended to focus on features that are most likely to positively impact on student behavior.

Though visualizing data is a vital consideration externally, a sequence of datamining processes is at the core of the framework. Data mining enables discovery of hidden information from many patterns of relationships that exist within these large, often noisy and messy datasets [19]. Eventually, these outer visualization and inner mechanisms of data mining aid human understanding and support learner's performance improvements [20].

The main research objective of this study is to build a dashboard plug-in that will focus on predicting the student's final academic performance. Various datamining approaches will be used along with visualization techniques to project online behavioral patterns. The LAD will provide students and educators a platform to evaluate the student's online behavior which will influence the students to perform better especially for those students who are at-risk of failing a course.

The remainder of this study is structured as follows. In section II, the previous studies have been reviewed. Section III explains the classification techniques that will be used to build the early warning model and the evaluation metric. Section IV defines the experimental design and the description concerning the graphs that are to be used in this study, and the future evaluation study, followed by Section V with the conclusion.

II. EXISTING DASHBOARD REVIEWS

This section provides an overview of some dashboards being used by learners/students and educators/teachers. We examine the purpose of the dashboards, their limitations and challenges in their design.

LISSA (Learning dashboard for Insights and Support during Study Advice) focused on first-year students. The goal of this dashboard was to assist Study Advisers (SA) in helping their students to plan a more attainable program and not waste years on incorrect choices. Student grades and historical data were used for the LISSA analysis to create a prediction bar which covered the duration of the bachelor program in which the students were enrolled. Three different coding colors, namely green, orange and red, were used to represent potential course outcomes of success, mediocre pass and fail. The proposed dashboard focused on visualizing only the academic performance of the students along with that of their peers [3]. Further, this tool was useful to students only with proper guidance from their SA. An overview of their study progress was provided by SAs to students near the end of the semester which was not helpful to those students who were about to fail a course. The retrospective nature of LISSA meant that it lacked currency in giving student feedback and also required assistance from the SA to make meaningful use of it. The tool was merely used to support the advising session and predictive algorithm was not used for predicting the students results. Moreover, the tool did not provide the students with actionable insights due to unchangeable data.

Arnold and Pistilli [11] from Purdue University created Course Signals, a dashboard for both students and faculty staff to track student progress. The data comprised student performance in the course, student engagement in terms of interaction with the LMS, academic history such as GPA, and demographics. The output is presented using a traffic light that shows three risk levels, high (red), medium (yellow) and low (green) and is sent to students via email, text, or the LMS. The tool determines whether a student is at risk of failing or withdrawing from a course as early as the second week of the semester. Use of the tool for different courses had shown great promise for first- and second-year students and in terms of the overall student retention at the university. The signals are updated only when the instructor runs the application. Moreover, the tool did not provide any direct insight into the specific causes behind a student being at risk or failure, rather it generates a prediction that identifies the level of risk, thus making it difficult to recommend a specific remediation.

Kuzilek et al. [12] from Open University, UK, proposed to predict at-risk students as early as possible. Therefore, a series of students who may be at risk of not submitting their next assignment is sent weekly to the course coordinators. A dashboard with two views was implemented to present (1) the predictions and a course overview and, (2) a view of individual students. Demographic (static) data and student interactions with the Virtual Learning Environment (VLE) system were used for the analysis. The tool again did not provide any direct insight to the students. Also, the prediction results are not available on the dashboard, rather they are sent to the course team via email. An evaluation of the efficacy of the tool was not done.

A tool for accessing and analyzing students' behavior data in learning systems, called AAT (Academic Analytics Tool), [13] was developed for course designers to get feedback about how students use a course and their learning processes. AAT allows users to specify the data they are interested in analyzing for selected courses. The learning objects that AAT used for the analysis are course materials, forum postings, quizzes, and video and audio files, but the data must be stored in a database accessible via SQL. The data are analyzed for the course revision process. However, this tool did not provide any graphical visualization of the students' behavior and no usability testing of the tool was performed.

Designed by Lonn and Teasley [22], Student Explorer is an early warning system (EWS) for academic advisers to support them in providing academic advice to at-risk students who are underperforming in their classes. Information from the LMS gradebook and assignments tool is used to measure students' formative performance. Three E3 categorizations are applied with labels to denote the most appropriate action by the academic adviser. The results indicated that the advisors mostly used the EWS during their meetings with students, despite it being designed as a tool to identify students who may be struggling and provide them information to prepare for such meetings.

The NTU student dashboard was implemented by Nottingham Trent University (NTU) to help students engage more effectively with their courses and found engagement to be a stronger predictor of success than the student background characteristics. Data collected by the system include student information, student access logs to university facilities, library loans, student interactions with online learning systems, electronic submissions of coursework and assignments, class attendance and tutor comments. The dashboard had helped to build better relations between students and the tutors, and some learners reported that seeing details of their own engagement encouraged them to continue to engage in their studies [14]. However, no information was provided regarding the effect of the dashboard on student learning behavior, skill levels or learning outcomes.

Gutiérrez et al. [16] discussed the development of the Learning Analytics Dashboard for Advisers (LADA) that supports academic advisers in gathering a semester plan for students based on their compiled history. Data from student grades, course credits, list of courses, and courses booked by a student were used in the analysis. The dashboard also includes a prediction of the chance of success for students and information regarding the quality of the predictions along with the graphs. An evaluation of the dashboard was undertaken via a pilot study in which the study advisers provided recommendations of courses for the students in the forthcoming semester. The results showed that the tool was found to be useful for laymen as compared to experts.

Among the reviewed papers, the main resources that have been used to gather data is students log data from Learning Management System (LMS) which includes student grades, demographic information, attendance and GPA.

Visualization of data plays a significant role in dashboard. The most widely used representations among the reviewed papers were bar graph where rich information

regarding students' activities are effectively delivered and line graph for delivering weekly login trends. Hence, these two graphs were widely used. This was followed by scatter plot for comparing individual students with other students and sociogram for presenting online communication and networking. Traffic signal was used to augment grades.

Most of the developed tools either concentrated on visualizing the students' predicted results or their online behavior but predicting the students' final academic performance along with visualization helps to identify the at-risk students thus making a recommendation for remediation easy. This paper concentrates on both data mining techniques along with visualizations to assist particularly at-risk students to improve their performances in their courses.

III. METHODOLOGY

There are two technological components to the proposed dashboard. The first is a visualization tool which provides students with a snapshot of how they are tracking with their levels of engagement and, where their engagement levels are situated in respect to the rest of the class. The second component makes use of machine learning algorithms which extract patterns from historic learning and engagement patterns of previous student cohorts in order to generate models that can predict outcomes early on in a course "Fig. 1".

The process will start with data gathering step where data will be collected from different learners' activities when they interact with various sources such as Moodle, Library and SMS. After gathering the data from various sources, the subsequent step is to clean the data of inconsistencies, errors and noise before, the data is mined. Data cleaning tries to fill in missing values and correct irregularities in the data. Data transformation will be done next to change the raw data into a specified format

according to the need of the model. Hence after cleaning the data, it becomes possible to do prediction using the selected features. The objective of the prediction is to evaluate the unknown value of a variable that defines the student final academic performance.

A. Prediction Algorithm

Predicting the numerical value in data mining is difficult to carry out accurately hence, the students will be next categorized into three groups [23]. Students in the lower bound of "55" (≤ 55) will be considered high risk of failing the final exam, while those in the higher bound of "75" (≥ 75) will be a low risk. Students in the medium band of "55-75" (> 55 and < 75) will be considered as medium risk of failing the exam. The prediction of final academic performance will be performed using classification technique, a procedure which classifies a given instance into a set of discrete categories [17]. The value of a (categorical) attribute (the class) depends on the values of other attributes (the predicting attributes). A model can be built by using a training dataset where all the attributes are known including the class attribute and verified via a test dataset where the class attribute is unknown.

Classification techniques is chosen because it has been widely used in educational data mining. For instance, it will be used to classify students into different categories by their final marks depending on the outcome of activities they performed in a learning system [21]. There are a vast variety of classifiers in the literature but choosing the best classifier is not a simple task because they differ in numerous aspects such as learning rate, robustness, and amount of data required for training. Such a classification method can be remarkably important in supporting education and learning. To classify the students into three different categories, following machine learning algorithms will be used; Random Forest classifier (RF), Naïve Bayes (NB), Logistic regression (LR), and k-Nearest Neighbours (kNN). Performance of various algorithms will be compared using

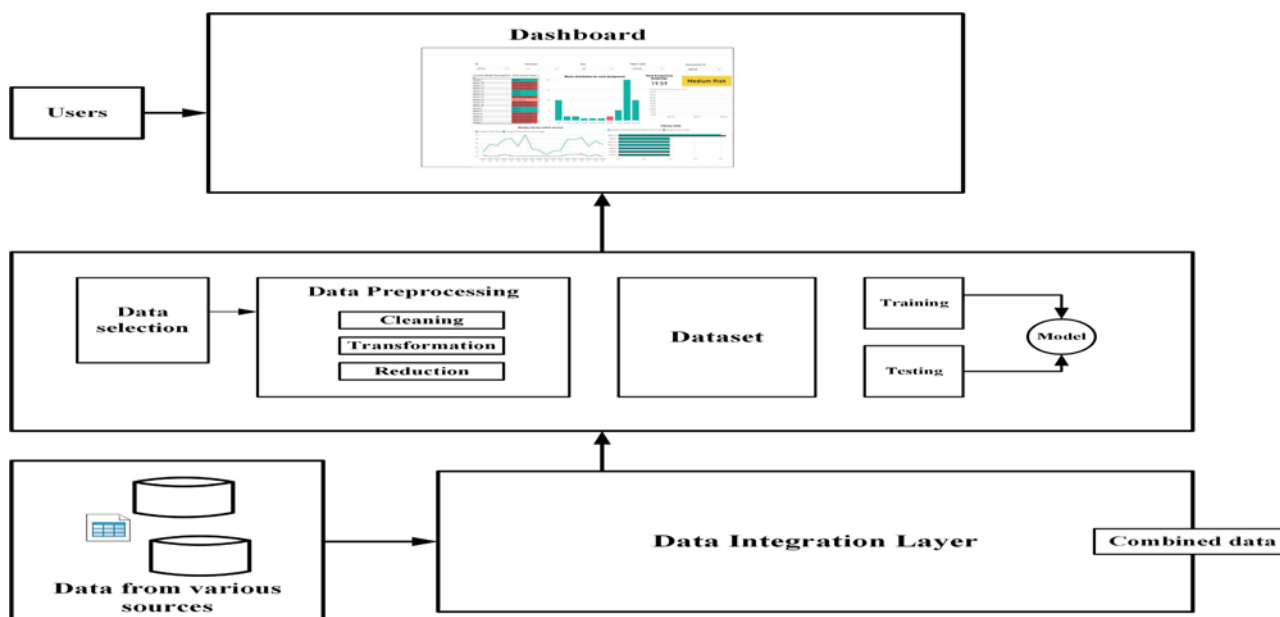


Fig. 1. System architecture

the hold-out and K-fold cross-validation methods and thus, the best performing algorithm will be selected. In an earlier study the authors have used the similar data mining technique to help an educator identify at-risk students, and found RF outperformed all the other classifiers [24]. Hence, in this study similar approaches will be used.

To evaluate a classifier's performance, the accuracy alone can be misleading if the dataset is highly imbalanced because a model can predict the value of the majority class for all predictions and attain a high overall classification accuracy [18]. Thus, correctly predicting the positive outcomes is not sufficient, and a predictive model should comprise a blend of both successful positive predictions and successful negative predictions. In such circumstances, it is more reliable to use F-measures which consider both the precision and recall and is defined as the harmonic mean of both the precision and recall.

$$F\text{-measures} = 2 * \frac{\text{Precision} * \text{Recall}}{\text{Precision} + \text{Recall}} \quad (1)$$

Precision is the fraction of relevant instances among the retrieved instances [6]

$$\text{Precision} = \frac{TP}{TP + FP} \quad (2)$$

Recall is the fraction of relevant instances retrieved over the total number of retrieved instances [6]

$$\text{Recall} = \frac{TP}{TP + FN} \quad (3)$$

where true positive (TP) and true negatives (TN) occurs when an object of a class 'yes' is predicted correctly as 'yes', and if the object of a class 'no' is wrongly predicted as 'yes' then it is called false positive (FP) and false negative (FN).

After predicting a student's final exam score the next step is to project these results along with user interactions alongside the institutional learning platform to the target users. This can be performed by developing a custom dashboard.

IV. DASHBOARD DESIGN

The dashboard will be implemented following a client-server architecture. On the client-side the Power BI tool will be used to develop the web application. Power BI combines connectivity to numerous data sources with a simple interface that makes dashboard creation fast and easily to distribute to numerous users. The dashboard has been first built on Power BI Desktop "(Fig. 2)". The deployment of the

developed dashboard is simplified by using Power BI service, which is built on Azure, a Microsoft's cloud computing infrastructure. Meanwhile, on the server-side, Python library scikit-learn [7] will be used to predict the academic risk of a student.

A. Dashboard prototype

Data privacy protections are one of the general requirements that need to be considered for LAD tools which manipulate and store personal data, as well as present it to different stakeholders [9] at varying degrees of aggregation. Study approval was obtained from the research ethics committee at the university to use students' data. A preliminary dashboard has been developed primarily using the student grades and library engagement data.

Library engagement data represents student actions such as the number of digital resources accessed by the students, number of physical visits to the library by students, and the number of physical contents borrowed from the library by the students. All these data were provided on weekly basis, followed by assessment scores of the students, which are derived from the institution's SMS.

The dashboard design was driven by the four phases of the LA process model defined by Verbert et al., [4]. According to the model, the dashboard should provide awareness to students of how their activity compares to that of their cohorts (phase one). This depends on visualization. Based on this information, students may reflect on their behavior (phase two). The third phase defined by is "sensemaking" where users gain new insights about what the data means. Finally, changes in performance and outcomes (phase four) of the students might occur.

The indicators on the dashboard have been depicted as simple line, bar, histogram and box plots. These types of indicators were used by other researchers in their educational dashboards due to the high interpretability they extend to the stake holders. Few [9] have recommended to use simple visuals to convey information to the users.

An individual student's activity or performance is compared with the average of the class, rather than showing every student individually on the visualization, thus preserving privacy. Comparing the students' performance with the class average has been one of the most extensively used approaches in learning analytics dashboards and visualizations [10]. This hierarchical layout will help students to get a better overview about their current learning pattern and status.

To efficiently communicate the prediction of the predictive models to the users, a traffic lights (signified as E in Fig 2) that corresponds to values of the prediction results along with descriptive words that shows three risk levels, "high risk" (red), "medium risk" (orange) and "low risk" (green) are presented to display the predicted final academic performance of students.

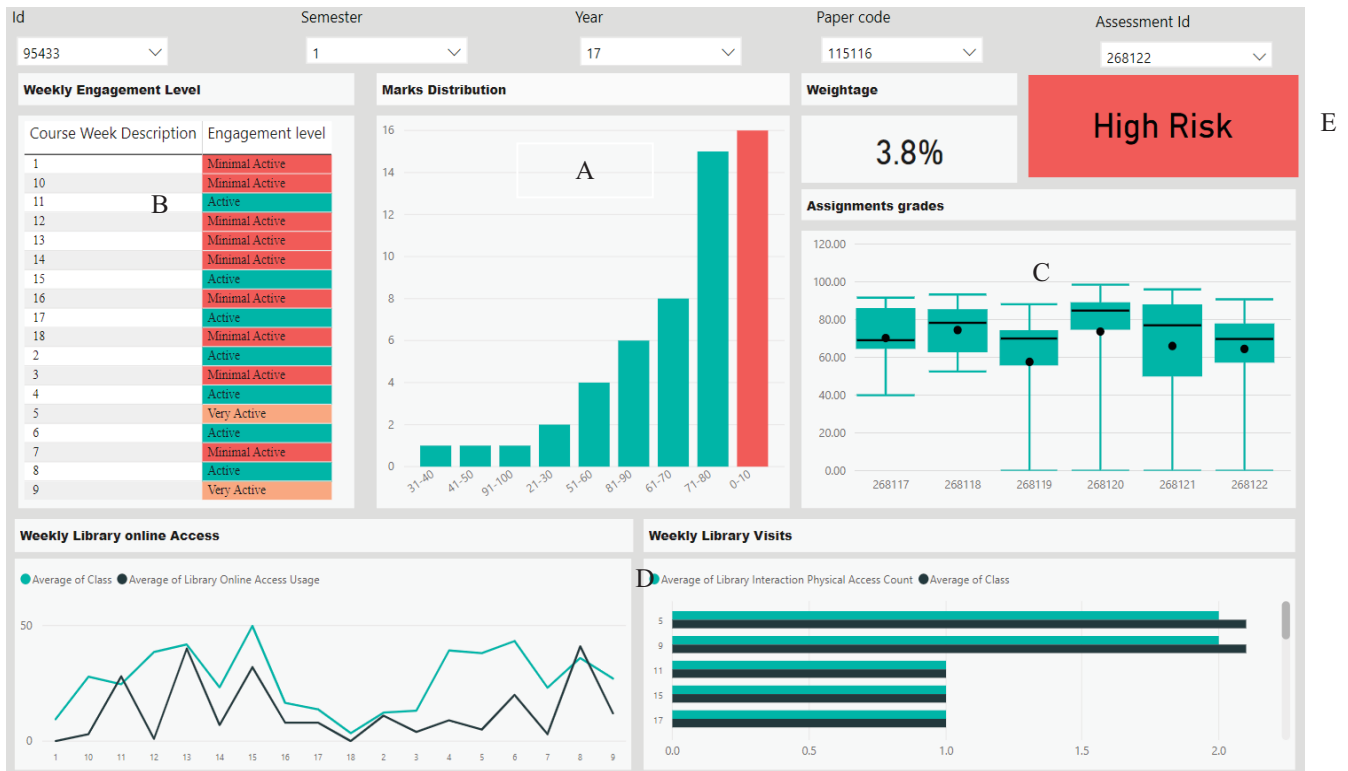


Fig. 2. A learning analytics dashboard designed for students

The histogram on the dashboard (signified as A in Fig 2) represents the student’s position (highlighted by red color) in relation to grades of their peers for every assignment.

The line and the bar graph (signified as D in Fig 2) show weekly library online access patterns of the students and their physical visits to the library. The week-by-week engagement graph of library data helps a student to gain an overall picture of their engagement compared with the rest of the class. The total percentage of credits achieved by the students in the assignments is also presented on the dashboard. Finally, table (signified as B in Fig 2) represents the “Activity behavior” of the students with the library resources and is categorized as “active” (orange), “very active” (green) and “minimal active” (red) and shows their weekly distribution over an entire term. A student is defined to be “active” if the engagement with the library resources is more than three per week. A student is defined to be “very active” if the average engagement with the library resources is more than four a week. The box and the whisker plot (signified as C in Fig 2) further enable a detailed view of the grades of the students for various assignments. The section inside the rectangle shows the median and “whiskers” above and below the box demonstrates the positions of the minimum and maximum recorded scores.

On the top of the dashboard is the filtering menu that displays the general data about the course, semester and year that helps the user to choose which information the indicator should present. The grades and the weightage will change based on the “Assessment id” and the “Paper code” chosen by the students. But not all filters can be used for each context because the usage of the library resources will not change irrespective of the filter chosen by the students.

A. Evaluation Study

It is noteworthy that most of the evaluations in the reviewed papers have addressed usability, usefulness, or user satisfaction rates with the intention of gathering feedback to improve the dashboards them self. Only very few studies actually looked at the impact of the dashboard on student performance. Most of the studies were either in the prototype state or the dashboard had not been implemented in an authentic educational situation.

We intend to study the impact of the dashboard on student learning outcomes. This will be conducted next as a pilot study by implementing the dashboard to courses in a real-world setting. Furthermore, to validate the effectiveness of the dashboard, it is intended to conduct experiment between two groups of students by dividing the students into treatment and control groups. The participants will be assigned randomly to the groups. This will help to attribute the difference in results between the two groups and thus it will ensure the validity of the dashboard. The students in the treatment group will be invited to use the dashboard by providing a link to it through Moodle. Since the visualizations will be displaying students online learning patterns on the dashboard, which will be drawn after a sufficient number of log files have been accumulated, students will be allowed to access and use the dashboard a few weeks after the semester begins. The students in the control group will not be invited to use the dashboard. The statistical t-test will be conducted to verify the presence of a statistically significant difference between the dashboard treatment group and control group. The final performance score of the students will be used as the post test score. The mean score of both the treatment group and control group will be compared to find out the difference. This will be conducted as a future study.

V. CONCLUSION AND FUTURE SCOPE

The extensive blending of digital technology into higher education (HE) impacts both teaching and learning practices and enables the tracking of most of the user interactions. These digital footprints can then be used to determine the academic progress of the students, predict students that are at-risk, and identify potential problems at an early stage with the purpose of tailoring educational delivery to each student's level of need. The present study focuses on developing a dashboard that capitalizes on online user activities alongside their educational data (from log files) to provide data-driven insights in student learning contexts, in order to maximize student learning outcomes. As an initial analysis the dashboard was developed using SMS and library data. The context of our future work will be broadened through integration of additional data from Moodle, a communication platform for visualizing the student's online behavior which can provide additional insights. Publishing materials and communication with students are the most frequent amongst used data for analysis in Moodle. The manner in which students interact with the dashboard is important because this will determine the effect it has on positively modifying their behavior. Hence, the impact of the dashboard on the student's outcome will be conducted as a pilot study by deploying the dashboard to the classes in a live setting. The evaluation will be done to find out whether students consider the dashboard useful in modifying their learning behavior and motivating them to maximize their learning outcomes.

VI. REFERENCES

- [1] R. S. Baker, "Stupid tutoring systems, intelligent humans," *International Journal of Artificial Intelligence in Education*, pp. 600-614, 2016.
- [2] G. Siemens and R. S. d Baker, "Learning analytics and educational data mining: towards communication and collaboration," in *Proceedings of the 2nd international conference on learning analytics and knowledge*, Vancouver, British Columbia, Canada, 2012.
- [3] S. Charleer, A. V. Moere, J. Klerkx, K. Verbert and T. De Laet, "Learning Analytics Dashboards to Support Adviser-Student Dialogue," *IEEE Transactions on Learning Technologies*, pp. 389-399, 2018.
- [4] K. Verbert, D. E. J. Klerkx, S. Govaerts and J. L. Santos, "Learning Analytics Dashboard Applications," *American Behavioral Scientist*, pp. 1500-1509, 2013.
- [5] K. Verbert, S. Govaerts, E. Duval, J. L. Santos, F. Assche, G. Parra and J. Klerkx, "Learning dashboards: an overview and future research opportunities," *Personal and Ubiquitous Computing*, pp. 1499-1514, 2014.
- [6] D. M. Powers, "Evaluation: from Precision, Recall and F-measure to ROC, Informedness, Markedness and Correlation," 2011.
- [7] F. V. G. Pedregosa, A. Gramfort, V. Michel, B. Thirion, O. Grisel, ... and J. Vanderplas, "Scikit-learn: Machine learning in Python," *Journal of machine learning research*, pp. 2825-2830, 2011.
- [8] D. Ifenthaler, "Learning Analytics: Benefits and Challenges for Higher Education," The council on Business and Society, 2015.
- [9] S. Few, *Information dashboard design*, O'Reilly, 2006.
- [10] L. Corrin and P. de Barba, "Exploring students' interpretation of feedback delivered through learning analytics dashboards," in *In Proceedings of the ascilite 2014 conference*, 2014.
- [11] K. E. Arnold and M. D. Pistilli, "Course Signals at Purdue: Using Learning Analytics to Increase Student Success," in *In Proceedings of the 2nd international conference on learning analytics and knowledge*, 2012.
- [12] J. Kuzilek, M. Hlosta, D. Herrmannova, Z. Zdrahal and A. Wolff, "OU Analyse: analysing at-risk students at The Open University," in *Learning Analytics and Knowledge conference*, 2015.
- [13] S. Graf, C. Ives, N. Rahman and A. Ferri, "AAT – A Tool for Accessing and Analysing Students' Behaviour Data in Learning Systems," in *In Proceedings of the 1st international conference on learning analytics and knowledge*, 2011.
- [14] N. Sclater, A. Peasgood and J. Mullan, "Learning Analytics in Higher Education: A review of UK and international," *Jisc*, 2016.
- [15] S. Malik, *Enterprise dashboards: design and best practices for IT*, John Wiley & Sons., 2005.
- [16] F. Gutiérrez, K. Seipp, X. Ochoa, K. Chiluita, T. De Laet and K. Verbert, "LADA: A learning analytics dashboard for academic advising," *Computers in Human Behavior.*, 2018.
- [17] C. Romero and S. Ventura, "Educational Data Mining: A Review of the State of the Art," *IEEE Transactions on Systems, Man, and Cybernetics*, pp. 601 - 618, 2010.
- [18] B. Jason, "How to estimate the performance of machine learning algorithms in Weka," 18 July 2016. [Online]. Available: <https://machinelearningmastery.com/estimate-performance-machine-learning-algorithms-weka/>.
- [19] Nisbet, Elder and Miner, *Handbook of Statistical Analysis & Data Mining Applications*, Elsevier, 2009.
- [20] S. K. Card, J. D. Mackinlay and &. Shneiderman, *Readings in information visualization: using vision to think*, San Francisco, CA, USA: Morgan Kaufmann, 1999.
- [21] M. Vahdat, "Learning analytics and educational data mining for inquirybased learning," The Netherlands, 2017.
- [22] S. Lonn and S. D. Teasley, "Student Explorer: A Tool for Supporting Academic Advising at Scale," in *In Proceedings of the first ACM conference on Learning@ scale conference*, Atlanta, 2014.
- [23] J. Vandamme, N. Meskens and J. Superby, "Predicting Academic Performance by Data Mining Methods," *Education Economics*, vol. 15, no. 4, pp. 405 -419, 2007.
- [24] G. Ramaswami, T. Susnjak, A. Mathrani, J. Lim and P. Garcia, "Using educational data mining techniques to increase the prediction accuracy of student academic performance," *Information and Learning Sciences*, vol. 120, no. 7/8, pp. 451 - 467, 2019.



Date: 19 June 2019

Dear Gomathy Suganya Ramaswami

Re: Ethics Notification - **NOR 18/71 - Learning Analytics: On Effectiveness of Dashboarding for Maximizing Students Outcomes**

Thank you for the above application that was considered by the Massey University Human Ethics Committee: **Human Ethics Northern Committee** at their meeting held on **Wednesday, 19 June, 2019**.

Approval is for three years. If this project has not been completed within three years from the date of this letter, reapproval must be requested.

If the nature, content, location, procedures or personnel of your approved application change, please advise the Secretary of the Committee.

Yours sincerely

Professor Craig Johnson
Chair, Human Ethics Chairs' Committee and Director (Research Ethics)



Date: 17 August 2020

Dear Gomathy Suganya Ramaswami

Re: Ethics Notification - **NOR 19/55 - Learning Analytics: On Effectiveness of Dashboarding for Maximizing Students Outcomes**

Thank you for the above application that was considered by the Massey University Human Ethics Committee: **Human Ethics Northern Committee** at their meeting held on **Friday, 14 August, 2020**.

Approval is for three years. If this project has not been completed within three years from the date of this letter, reapproval must be requested.

If the nature, content, location, procedures or personnel of your approved application change, please advise the Secretary of the Committee.

Yours sincerely

Professor Craig Johnson
Chair, Human Ethics Chairs' Committee and Director (Research Ethics)