

A COMPARATIVE ANALYSIS OF SELECTED STUDIES IN STUDENT PERFORMANCE PREDICTION

Michael Mogessie Ashenafi

Department of Information Engineering and Computer Science, University of Trento,
Italy

ABSTRACT

Predicting student success has long been an interest of institutions of higher education as well as organisations responsible for preparing high-stake, standardised tests administered at national and international levels. This study discusses how performance prediction studies have evolved from those that use demographic data and high school grades to predict success in college to those that utilise sophisticated data collected in non-traditional educational platforms to predict end-of-course performance and to those that show how student progress can be tracked in a continuous manner. A total of 56 studies published since the nineties are discussed. Views on strengths and weaknesses as well as observed opportunities for improvement are presented. The consistently high results reported in many of the studies shall convince the reader that automated solutions to the problem of predicting student progress and performance can either be tailored for specific settings or can be adopted from similar settings in which they have been utilized successfully. A recommendation on how to build upon recent success is provided.

KEYWORDS

Student performance prediction, prediction models, predictive analytics, educational data mining

1. INTRODUCTION

Assessment in education is usually carried out to achieve one of two goals. The more common of the two is summative assessment, which is intended to measure the extent to which students have achieved pre-specified learning goals and is usually carried out at specific intervals. Midterms and final exams are the most common examples of such type of assessment. Formative assessment, on the other hand, is carried out with the goal of guiding and supervising students to achieve expected learning goals. It allows students to continuously measure their progress and to identify their strengths and weaknesses in grasping a subject matter.

Regardless of the stark differences between the goals they set out to achieve, the two forms of assessment share the general objective of measuring student performance. Among other factors, the effectiveness of this measurement largely depends on its design and timeliness.

The effectiveness of standardised and high-stake summative assessment practices in measuring student performance is open to debate. The intended purpose of this study, however, is not to shed more light on this topic but to demonstrate how researchers have turned to data analysis and machine learning techniques to predict outcomes of widely used assessment practices such as midterms and final exams.

Traditional assessment tasks are usually prohibitive by nature as they require substantial amount of time to design and administer. For this reason, it is rare to see summative assessment tests being continuously administered throughout a course. The implied burden of continuous

summative assessment does not entail that it is ineffective. It, nonetheless, makes it unpopular amongst those who are proponents of formative assessment and other student-centric practices.

The major goal of prediction systems is to provide timely information that would help stakeholders to either make the required preparations in to handle predicted events well or to make the necessary adjustments to avert harmful events. Performance prediction is no different in that it is carried out with the intention of providing students and teachers information that can be used to measure progress and to identify those at risk of failure well ahead of time so that appropriate measures could be taken both on parts of the teacher and the student to avoid such risk. For this reason, timeliness shall be regarded as a necessary albeit a hardly sufficient condition for the effectiveness of performance prediction and other prediction models in general.

This study investigates a selected number of studies from the student performance literature available to-date to determine whether performance prediction practices have evolved significantly over the years. It examines parameters that have been of interest in making such prediction. It also discusses the role that advances in computer science and related disciplines have played and continue to do so in improving research in this area. The study highlights the importance of using prediction models to provide timely feedback and how this can be achieved using continuous prediction of student progress.

Research in performance prediction has led to a significant number of publications over the past forty years. An effective review of such literature should have several themes as its focus. Initial investigation of the literature has revealed topics and educational settings that have been of specific interest in this research area. Most studies have been conducted in higher education settings and apply one or more machine learning algorithms to build performance prediction models. For this reason, it was decided to limit the focus to the review of studies in higher education that applied machine learning techniques to student performance data collected in traditional or online learning environments, or both.

The phrase 'predict student performance' was used to carry out a literature search on IEE Explore, ACM digital library and Google Scholar databases. Initial filtering was applied to the top hundred results from each search based on the titles of the studies. A further selection among the remaining 131 studies excluded studies that did not specifically address student performance prediction or were not conducted in a higher education setting. Fifty-six studies were included in the final list of studies to be reviewed. Of these, 10 provided either unclear or little information to be included in the statistical summaries provided. Although a comparison of the studies according to several qualitative parameters also covers these 10 studies, the charts presented here use data from the remaining 46 studies.

The next section presents an analysis of how the sources of prediction data have evolved since the early performance prediction studies. Several dimensions were identified in the process of comparing the studies. Many of the studies are conducted in undergraduate-level computer science or engineering courses. Details of this observation are provided in section 3. In section 4, whether predictions focused on a simple outcome such as a pass or fail, or more complex outcomes such as specific marks or grades are used to compare the studies. One-off versus continuous predictions made by the studies are addressed in section 5. Investigation of the studies to determine what parameters may prove to be good predictors of student performance and what factors may lead to better results is presented in section 6. Section 7 and 8 focus on the array of prediction algorithms used in the studies and on studies that utilised large amounts of data from Massive Open Online Courses (MOOCs), respectively.

Common evaluation metrics may not provide accurate information about the performance of a prediction model that serves as a tool of early intervention. This is because the major goal of such a model is to avoid the predicted outcome. A brief note on this concern is provided in section 9. The final section provides a discussion and a few recommendations for the continued improvement and institutional-level adoption of student performance prediction.

2. EVOLUTION OF DATA SOURCES IN STUDENT PERFORMANCE PREDICTION

Revisiting earlier studies regarding student performance prediction reveals that they were born out of the debate whether standardized tests could predict student success in later years of higher education ([1], [2], [3]).

Consequently, earlier research focused on determining whether grades standardised test results that were obtained at earlier levels of education could predict well student success in later years of higher education ([4], [5], [6], [7]). Recent studies that still consider student performance at the lower levels of education incorporate additional parameters obtained from more recent achievements of students such as freshman and sophomore years as well ([8], [9], [10], [11]).

Most performance prediction studies have now become more refined in the parameters they utilise, the statistical measures they employ and the outcomes they predict. Demographic data, performance on take-home assignments, projects, and other tasks and activity on online learning and assessment platforms have all been used to build end-of-course or end-of-programme performance prediction models.

The advent of the Internet and the World Wide Web (WWW) have made it possible to construct online interaction environments, paving the way for the collection of student data in an unparalleled manner. Accordingly, recent research in performance prediction has augmented traditional parameters with student activity data obtained from online educational environments. This phenomenon has helped ease the transition from prediction of a binary outcome such as a pass or a fail to prediction of more sophisticated and fine-grained outcomes such as actual grades and continuous, numerical scores.

Examples of studies that utilised information from students' activities in online learning and assessment platforms in predicting performance include ([12], [13], [14], [15], [16], [17], [18]). Figures 1 and 2 show the studies reviewed here categorised according to the types and variety of the data sources they used. Many studies used multiple data sources. The chart in figure 1, therefore, shows a higher number than the actual 46 studies.

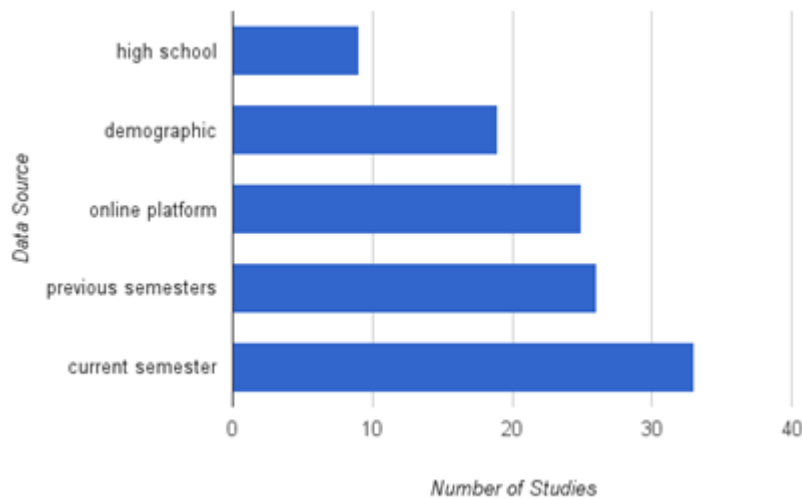


Figure 1: Studies by source of prediction data (N=46)

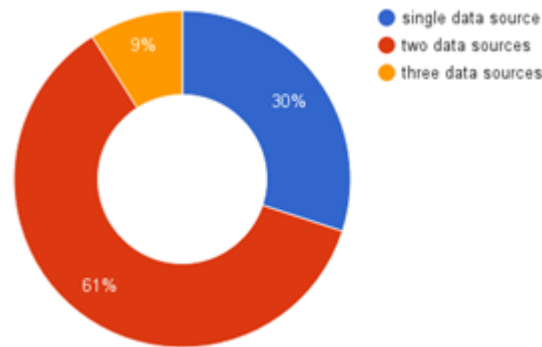


Figure 2: Studies by the number of data sources used (N=46)

3. COURSE LEVELS AND SUBJECT AREAS

Many the studies that reported the subject areas they were conducted in used data from courses administered as part of either computer science or engineering programmes at the undergraduate level. Of these, many focused on predicting performance of freshman and second year students enrolled in introductory level courses.

Reasons behind the lack of similar studies in other subject areas or at different course levels remain an assumption. Two partial explanations may be that utilising machine learning algorithms is an immediate advantage that researchers in the field of computer science have. It could also be that performance prediction has more impact when conducted at earlier years of college education. Figures 3 and 4 show the studies categorised according two subject areas and course levels.

4. GRANULARITY OF PERFORMANCE PREDICTION – OVERALL SUCCESS VERSUS SPECIFIC OUTCOME

Apart from the techniques they use, studies seeking to predict student performance differ in how they interpret performance. A generic approach is to predict an outcome that is binary in nature. Some of the studies reviewed follow this approach to predict overall outcome such as passing or

failing a course or even forecasting successful completion of college as marked by graduation ([19], [20], [21], [22], [23]). Others have taken a further step in predicting the classification of the degree or achievement ([24], [25]).

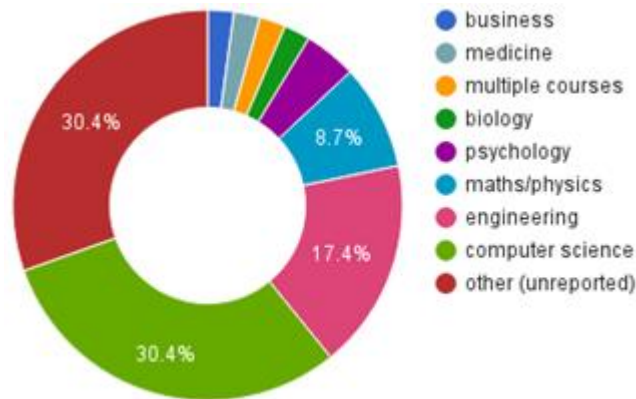


Figure 3: Studies by subject area (N=46)

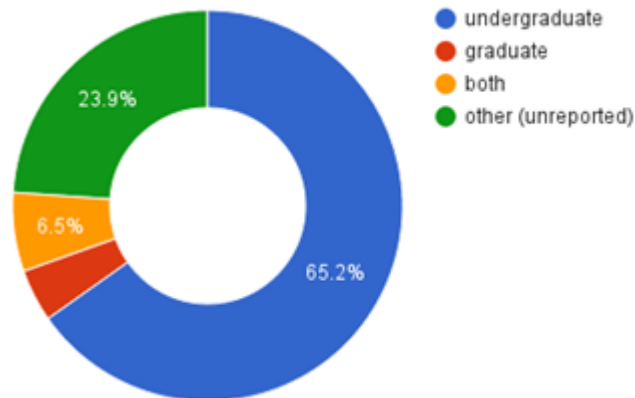


Figure 4: Studies by course level (N=46)

A performance prediction spectrum could be inferred from these studies, with prediction of a binary outcome, such as a pass or fail, placed on one end of the spectrum and prediction of a specific numerical score, on the other. Within these two extents lie prediction of a range within which a score would fall and prediction of numerical or letter grades.

Figure 5 shows the studies grouped according to the outcomes they predicted. Some of these studies predicted overall success as well as specific outcomes. Predicting overall success in a timely manner serves as a mechanism of early intervention. Timely prediction of a range or scores or grades has, however, a more powerful formative value as it provides granular information about specific performance categories of students ([26]).

In recent years, prediction of student performance has become more fine-grained and sophisticated. Researchers now seek to predict actual scores for tests and assignments as well as final scores and grades for an entire course. In its simplest form, effective prediction of such outcomes has a two-part requirement. The type and amount of student data that is to be collected forms one part and the data analysis and choice of prediction techniques form the other.

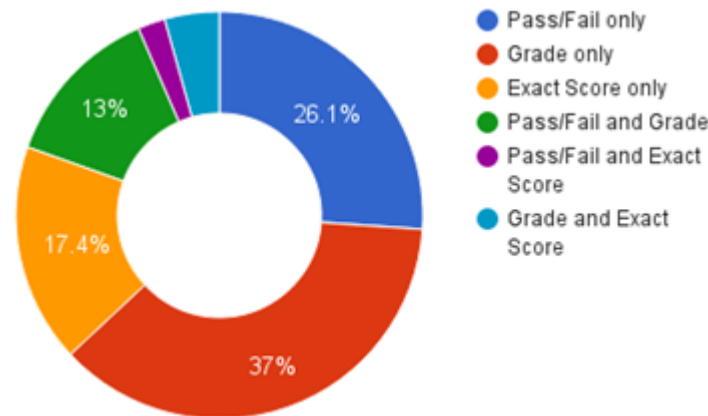


Figure 5: Studies by type of prediction (N=46)

5. ONE-OFF VERSUS CONTINUOUS PREDICTION – THE CASE OF SUMMATIVE AND FORMATIVE PREDICTION MODELS

Performance prediction models provide several advantages for both students and teachers. In summative assessment environments, prediction only provides information about end-of-course performance or overall performance in later years of higher education. In contrast, student-centred performance prediction aims to continuously provide information on the progress of students. Such information is helpful for students because, if properly utilised, it helps them identify their strengths and weaknesses in certain topics in a course and gives them the opportunity to act accordingly. It also helps the teacher measure the overall progress of the class and identify those students that may require special supervision. Borrowing from assessment terminology, student performance prediction can therefore be categorised into formative and summative.

The parameters that are chosen to act as predictors of student performance essentially determine the type of prediction that is to be made. Some predictor attributes, once obtained, are highly unlikely to evolve over time. Hence, it is claimed that although they may provide important information about expected performance when used together with other dynamic variables, they hardly contribute to measuring student progress when used separately to build prediction models.

To demonstrate this, a scenario where student success at the end of the first year of college is predicted when the student is about to enrol in the first-year courses is considered. A common approach for building such prediction models is to train a machine-learning algorithm with student performance data collected from the final years of high school combined with students' demographic data. It has been shown such data are indeed good predictors of success at a critical point such as the end of the first year of college studies ([6], [27], [28], [29], [30], [31], [32], [33]). However, much of the demographic data seldom changes and academic performance history from earlier years never does. For this reason, if a student's performance is predicted halfway through the first year using the same model, the predictions remain unchanged. Consequently, the use of such static parameters alone only allows making one-off predictions.

When combined with data that changes over time, demographic and previous academic performance data provide more information and can be used to measure progress. While static data provide information about a student's background, dynamic data such as the number of assignments completed to-date provide progress information. Together, static and dynamic data constitute a student profile.

There is no specific combination of static and dynamic student data that should apply to all prediction models. However, it should be noted that the higher the proportion of static data the less powerful the model would be in predicting progress. Therefore, studies that investigate whether it is possible to provide continuous performance predictions to students should weigh the worth of including static parameters in the prediction process.

One way of predicting progress is through predicting performance at specific intervals within a course. For instance, [34] show how to predict performance on specific course modules.

In general, experiments that aim to predict student progress could be set up in a way that predictions made at a specific point during the course utilise data from previously enrolled students up to the specific point of prediction. This could provide information on how students from previous batches with similar progress levels performed at the end of the course. In experiments that utilise machine learning algorithms, this implies training the models with a portion of the available data set that covers the period from the start of the course to the specific point of prediction only, and not on the entire data set. [26] Demonstrate this approach using peer-assessment data.

Online educational platforms provide a unique advantage of collecting dynamic data and are the perfect environment to build progress prediction models upon. Studies that utilise such platforms to provide predictions at several intervals of a course include [35], [36], [15]. These studies serve as evidence that static data can be done away with when predicting progress, without affecting prediction performance.

Because one-off predictions usually focus on predicting end-of-course performance, they hardly provide information about progress, which is important for students if they are to adjust accordingly. In this manner, such predictions may be considered as having only summative value. Nonetheless, about half of the studies that made one-off predictions (N=19) could be transformed to make continuous predictions because the prediction models could be applied to incomplete data obtained at several intervals during the course.

Figure 6 shows the studies categorized according to whether they make one-off (N=21) or continuous predictions (N=6) and those that make one-off predictions but could be transformed to provide continuous predictions (N=19).

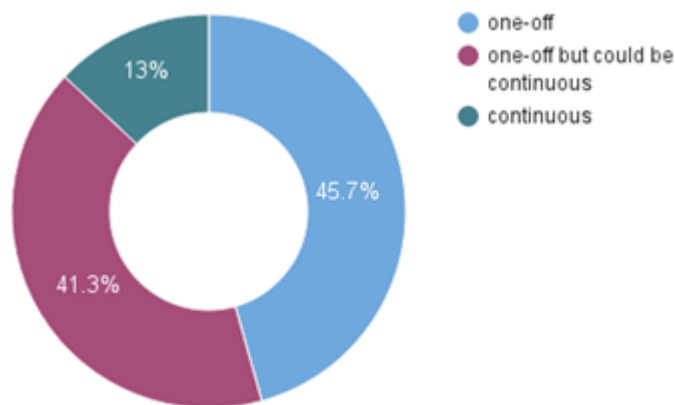


Figure 6: Studies by frequency of prediction (N=46)

6. WHAT ARE GOOD PREDICTORS OF STUDENT SUCCESS AND HOW GOOD ARE THE PREDICTIONS?

Despite the abundance of studies in performance prediction and the parameters used in making predictions, it is difficult to single out parameters as outstanding predictors that have consistently significant values across the studies reviewed. Two immediate reasons for this may be the variation in the setup of experiments and the prediction algorithms used. Other factors include the amount of data and the course or discipline within which the studies were conducted.

Nonetheless, studies that used a higher number of predictor parameters and a larger number of student records reported higher results. Because online learning and assessment environments simplify the collection of student activity data, and hence allow inclusion of parameters about such activity in the prediction process, studies that used online educational platforms appeared to report significantly better results.

Similarly, predictors that were built using data from a large cohort of students or utilised data collected over long periods of time, had high performance accuracy. It should, however, be noted that accuracy may not be the right choice of performance evaluation technique, especially when prediction is essentially a classification task. In such scenarios, other metrics such as precision, recall, F1 scores and False Positive Rates could be used as they provide more representative and realistic information about how well a prediction model performs. This is especially true when a prediction model is evaluated according to how many at-risk students it identifies and how many it fails to do so.

Most the studies used previous academic performance data (N=21), many in combination with other data such as online activity logs. Fourteen of the studies used demographic data in making predictions. A comparison of the data sources and the reported performance results revealed that demographic data are more effective when used in combination with two or more other data sources such as online activity logs and previous academic performance data. Indeed, those studies that used solely demographic data or in combination with only one other data source reported moderate performance levels.

In contrast, studies that included partial marks, mid-term results and assignment scores reported higher results. The highest results were reported by studies which included three or more of these parameters in their predictions.

7. PREDICTION TECHNIQUES AND ALGORITHMS

Many the studies reviewed follow the approach of applying a range of machine learning algorithms to their data set and choosing the algorithm that reports the highest level of performance. The most common of these algorithms are Linear Regression, Neural Networks, Support Vector Machines, Naïve Bayes Classifier, and Decision Trees. A few studies use a combination of classifiers for improved predictions ([37], [24], [38]).

Other studies that follow less common approaches include those that use Markov Networks ([39]), Collaborative Multi-Regression models ([40]), smartphone data to investigate the correlation between students' social and study behaviour and academic performance ([41]) and those that perform Sentiment Analysis of discussion form posts in MOOCs ([42]). Yet, some studies discuss algorithms developed for the sole purpose of student performance prediction ([43], [44]).

The prediction algorithms used depend on the type of outcome being predicted. When predicting a binary outcome such as pass or fail, classification algorithms or algorithms that predict an outcome with a certain probability, such as Naïve Bayes Classifier, Decision Trees and Neural Nets are common. When predicting numerical scores, algorithms such as Linear Regression, Support Vector Machines may be applied.

Figure 7 shows the studies categorised by the type of prediction algorithm they used. Those that applied more than one algorithm are grouped under the ‘multiple’ category.

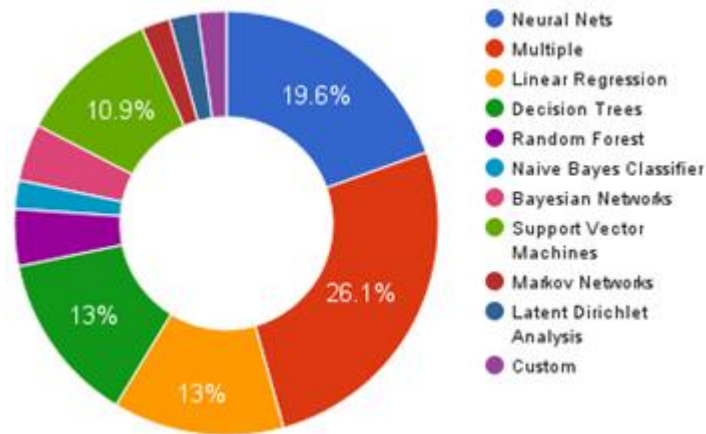


Figure 7: Studies by type of prediction algorithm

8. PREDICTIONS IN MASSIVE OPEN ONLINE COURSES (MOOC)

One of the factors that may weaken the appeal of MOOCs as complementary to, or even replacements of, traditional course administration practices is the fact that they are plagued by high attrition rates ([45]). It is, hence, not surprising that most student performance prediction studies that involve MOOCs have focused on predicting student dropout.

Because MOOCs are inherently tied to web-based platforms, the task of collecting data regarding student activity is only as challenging as building the platforms themselves. MOOCs allow gathering of immense amounts of data from students that do not necessarily attend courses from the same geographic region. Consequently, the diversity and size of the data collected by such platforms is unparalleled.

The impact of the unique advantage provided by using online course administration systems is immediately apparent in the performance levels of prediction models that are built on such data. Although the outcomes they predict are less complex when compared to models that predict grades or final scores, predictors in MOOCs consistently perform better than their counterparts built on top of traditional educational settings.

The abundance of data collected by such systems is not limited to improving prediction performance. In fact, many MOOC-related studies set out to explore other pedagogical aspects such as discovery of factors that determine student success in courses administered in a MOOC format. Among such factors are patterns captured from participation in online forums such as the length of posts in online discussion forums and whether a student is the creator of a discussion thread ([46]).

Other activities that may be used as features for building MOOC dropout prediction models include number of video lecture downloads, number of completed quizzes, number of completed tasks, click-stream data, the amount of time spent on course modules and the number of days students are active ([47], [25], [36], [35], [47], [48], [49], [50]).

9. THE POTENTIAL PARADOX OF AN EFFECTIVE PERFORMANCE PREDICTION SYSTEM

A straightforward approach to testing the performance of models that predict favourable outcomes is to compare such predictions with actual outcomes. Such an approach, however, may not apply to systems that predict unfavourable outcomes or even catastrophic events as the very essence of these predictive systems is rooted in their ability to provide information to those responsible for preventing the predicted events.

If common evaluation techniques are used to measure the performance of prediction systems that help avert unfavourable outcomes, evaluation results will lead to the conclusion that these prediction models do not perform well. By the same argument, a conclusion that these models are good predictors implies bearing the consequences of unfavourable outcomes. Although this paradoxical nature of evaluating prediction models is somewhat non-existent in models that predict natural catastrophes, it is still evident in models that provide timely prediction of student performance.

Concretely, a question is raised about how to test the performance of a model that predicts an outcome in a timely manner if the very performance of the models is rooted in its ability to provide information that avoids the predicted outcome. In the case of student performance prediction, timely prediction helps identify students who are at risk of failing. If that information is acted upon properly, the undesired outcome, failing a course is averted. Classical evaluation techniques such as prediction accuracy would provide the misleading conclusion that the model does not perform well, when it indeed does.

The performance of such non-traditional prediction models can be tested in two ways. An approach that is common to many well-designed research experiments is to divide subjects into experimental and control groups. In a class of students divided into such two groups, continuous and timely prediction will be provided to both groups but only those in the experimental group will be supervised according to the information provided by the prediction model. At the end of the course, analysing how many at-risk students in both groups, as predicted by the model, improved and eventually passed the course shall reveal the true performance of the prediction model. Traditional performance evaluation techniques could be used to evaluate the performance of the model on the control group. Statistical methods that measure how varied two sets of outcomes are can be used to judge how good the prediction model performs on the experimental group. Ideally, large differences between predictions and actual outcomes would imply that the prediction model has helped avoid the undesired outcome.

This approach, however, entails more risk for students in the control group as they may not benefit from the prediction and may indeed fail the course. Another approach that does not partition students into distinct groups involves periodic evaluation of only those students that the prediction model deems to be at-risk of failing. Although this method implies more work on part of the teacher, it provides invaluable information about the actual progress of the student. As noted in earlier sections, this approach is typical in formative assessment scenarios.

10. DISCUSSION AND CONCLUSION

Research in predicting student performance has been around for years. Initial investigations started as a way of seeking to establish whether standardised tests actually measure student performance in later years of education. Consequently, earlier studies leveraged statistical methods such as correlation coefficients to measure the relationship between such test results and exam results obtained during college education. The most common independent variables in such studies were composed of student demographic data and high school grades in several subjects. The dependent variables were often cumulative major or overall GPAs obtained in the first or last years of college.

Much of performance prediction research has focused on the disciplines of computer science and engineering. Although not far-fetched, the argument that researchers in these disciplines have at their disposal tools and know-how to build better mathematical prediction models may not explain the observation very well, especially when there is nothing to suggest that nature of the data used to build prediction models in these disciplines is essentially unique from those in other disciplines. It is hoped that the availability of computerised tools that provide an array of prediction algorithms serve as a motivation for researchers in other fields to carry out similar studies, which may support findings of current research.

Although earlier research sought to establish whether it was possible to make predictions that were binary in nature, such as predicting passing or failing a course, recent studies have managed to predict specific grades with high accuracy. In fact, some studies go as far as predicting specific scores, and with very small errors. In general, predicting outcome becomes more and more difficult as the number of possible outcomes grows. Hence, studies that predict actual scores have historically been less accurate than those that predict whether a student passes or fails a course. However, significant increase in performance is noted in score prediction models when the amount of data available for training the models is considerably large. However, this does not necessarily imply that prediction accuracy can be improved without bound by introducing more and more data. It is possible that after a certain point, prediction performance may not improve despite that introduction of new data. This is because although size of the data plays an important role, the prediction parameters and the algorithm in use are just as important. After all, it is learning algorithms that use many parameters – algorithms with low bias – and large training data, which provide low variance, that constitute a good prediction model.

Whether prediction is intended for summative or formative purposes highly determines the approach to building the model. While educational policy makers may be more interested in determining whether national and standardised tests are good predictors of long-term student performance, teachers benefit more from information that provides insight into student progress. Students benefit from timely information about their progress as well because it gives them the opportunity to act accordingly. The parameters that are chosen to build the model should reflect this notion of progress as well.

A plot of the number of publications per year for the studies reviewed indicated that research in performance prediction had progressed at a slow pace since the early nineties, with typically none to three studies published every year. However, it was observed that this has changed recently due to a sharp increase in the number of publications, with 41 studies published since 2010 alone. As evidenced in this study, this sharp increase goes hand in hand with the wide adoption by higher education institutions of online learning and assessment platforms, the most popular of which are Massive Open Online Courses (MOOCs).

The vast amount of educational data that is continuously generated in online learning environments has enabled the building of high performance and sophisticated prediction models. Performance levels of prediction models discussed in recent studies, most of which use data from online educational platforms, are consistently high. Nevertheless, more informative performance evaluation metrics shall provide the true performance levels of prediction models. When considering student performance prediction models, misclassifying students who are at risk as 'not at risk' has more consequences than if the reverse was true. Use of more robust evaluation metrics that capture such information is therefore in order.

Performance prediction algorithms provide more information about predictor variables than just performance statistics. The reason behind choosing a specific classification algorithm should not only be due to its better performance than other classification algorithms provided by the software utilised. The question of why a specific machine learning technique or statistical measure is more suitable than others is seldom addressed by the studies reviewed. The answers to this question are very likely to reveal additional, yet crucial, information about the data set under consideration and may even indicate which variables have higher impact on the predicted outcome.

Unfortunately, analysis of prediction results and how they may support or contradict common belief regarding predictors of student performance has often been a task left for the reader to ponder in most studies that iterate through each algorithm available in the toolkit and choose the one that outperforms all the rest.

Yet, there are exemplary studies that demonstrate how performance prediction studies should be conducted, how prediction parameters and results should be analysed, how comparisons with other approaches should be presented and that provide elaborate explanations of expected outcomes and peculiarities of observations. Examples of such research are discussed in various recent works ([38], [39], [44], [41], [42], [49], [15]).

Despite some pitfalls that exist in the design of some prediction models, research in student performance prediction has arguably reached a stage that warrants its integration into practical educational settings, albeit with caution. Indeed, studies that discuss how good prediction models have been integrated into existing courses with success shall encourage others that are contemplating such an integration.

REFERENCES

- [1] Kellaghan, T., Madaus, G., and Airasian, P. (1982). *The effects of standardized testing*. Boston: Kluwer-Nijhoff Pub.
- [2] Crouse, J., and Trusheim, D. (1988). *The case against the SAT*. University of Chicago Press.
- [3] Fincher, C. (1990). Trends in the predictive validity of the scholastic aptitude test. *ETS Research Report Series*, 1990(1), i-19. <http://doi.org/10.1002/j.2333-8504.1990.tb01345.x>
- [4] Fowler, G. C., and Glorfeld, L. W. (1981). Predicting aptitude in introductory computing: A classification model. *AEDS Journal*, 14(2), 96-109. [doi:http://doi.org/10.1080/00011037.1981.11008293](http://doi.org/10.1080/00011037.1981.11008293)
- [5] Evans, G. E., and Simkin, M. G. (1989). What best predicts computer proficiency?. *Communications of the ACM*, 32(11), 1322-1327. <http://doi.org/10.1145/68814.68817>
- [6] Jamba-Joyner, L., and Klostermeyer, W. F. (2003). Predictors for success in a discrete math course. *ACM SIGCSE Bulletin*, 35(2), 66-69. <http://doi.org/10.1145/782941.782983>

- [7] Al-Hammadi, A. S., and Milne, R. H. (2004). A neuro-fuzzy classification approach to the assessment of student performance. In *Fuzzy Systems, 2004. Proceedings. 2004 IEEE International Conference on* (Vol. 2, pp. 837–841 vol.2).<http://doi.org/10.1109/FUZZY.2004.1375511>
- [8] Huang, S., and Fang, N. (2012). Work in progress: Early prediction of students' academic performance in an introductory engineering course through different mathematical modeling techniques. In *Frontiers in Education Conference (FIE), 2012* (pp. 1–2).<http://doi.org/10.1109/FIE.2012.6462242>
- [9] Kumar, S.A., and Vijayalakshmi, M. N. (2012). Mining of student academic evaluation records in higher education. In *Recent Advances in Computing and Software Systems (RACSS), 2012 International Conference on* (pp. 67–70).<http://doi.org/10.1109/RACSS.2012.6212699>
- [10] Tiwari, A. K., Pandey, A., Rohatgi, D., and Singh, A. K. (2013). Result prediction system for Multi-Tenant database of University. In *Advance Computing Conference (IACC), 2013 IEEE 3rd International* (pp. 1340–1344).<http://doi.org/10.1109/IAdCC.2013.6514423>
- [11] Arsad, P. M., Buniyamin, N., and Manan, J.-L. A. (2013). A neural network students' performance prediction model (NNSPPM). In *Smart Instrumentation, Measurement and Applications (ICSIMA), 2013 IEEE International Conference on* (pp. 1–5).<http://doi.org/10.1109/ICSIMA.2013.6717966>
- [12] Aguiar, E., Chawla, N. V, Brockman, J., Ambrose, G. A., and Goodrich, V. (2014). Engagement vs Performance: Using Electronic Portfolios to Predict First Semester Engineering Student Retention. In *Proceedings of the Fourth International Conference on Learning Analytics and Knowledge* (pp. 103–112). New York, NY, USA: ACM.<http://doi.org/10.1145/2567574.2567583>
- [13] Ashenafi, M. M., Riccardi, G., and Ronchetti, M. (2015). Predicting students' final exam scores from their course activities. In *Frontiers in Education Conference (FIE), 2015 IEEE* (pp. 1–9).<http://doi.org/10.1109/FIE.2015.7344081>
- [14] Simjanoska, M., Gusev, M., and Bogdanova, A. M. (2014). Intelligent modelling for predicting students' final grades. In *Information and Communication Technology, Electronics and Microelectronics (MIPRO), 2014 37th International Convention on* (pp. 1216–1221).<http://doi.org/10.1109/MIPRO.2014.6859753>
- [15] Koprinska, I., Stretton, J., and Yacef, K. (2015). Predicting Student Performance from Multiple Data Sources. In C. Conati, N. Heffernan, A. Mitrovic, and M. F. Verdejo (Eds.), *Artificial Intelligence in Education SE - 90* (Vol. 9112, pp. 678–681). Springer International Publishing.http://doi.org/10.1007/978-3-319-19773-9_90
- [16] Grivokostopoulou, F., Perikos, I., and Hatzilygeroudis, I. (2014). Utilizing semantic web technologies and data mining techniques to analyze students learning and predict final performance. In *Teaching, Assessment and Learning (TALE), 2014 International Conference on* (pp. 488–494).<http://doi.org/10.1109/TALE.2014.7062571>
- [17] Manhães, L. M. B., da Cruz, S. M. S., and Zimbrão, G. (2014). WAVE: An Architecture for Predicting Dropout in Undergraduate Courses Using EDM. In *Proceedings of the 29th Annual ACM Symposium on Applied Computing* (pp. 243–247). New York, NY, USA: ACM.<http://doi.org/10.1145/2554850.2555135>
- [18] Zhou, Q., Zheng, Y., and Mou, C. (2015). Predicting students' performance of an offline course from their online behaviors. In *Digital Information and Communication Technology and its Applications (DICTAP), 2015 Fifth International Conference on* (pp. 70–73).<http://doi.org/10.1109/DICTAP.2015.7113173>
- [19] Nghe, N. T., Janecek, P., and Haddawy, P. (2007). A comparative analysis of techniques for predicting academic performance. In *Frontiers in Education Conference - Global Engineering: Knowledge without Borders, Opportunities without Passports, 2007. FIE '07. 37th Annual* (pp. T2G–7–T2G–12).<http://doi.org/10.1109/FIE.2007.4417993>

- [20] Karamouzis, S. T., and Vrettos, A. (2009). Sensitivity Analysis of Neural Network Parameters for Identifying the Factors for College Student Success. In *Computer Science and Information Engineering, 2009 WRI World Congress on* (Vol. 5, pp. 671–675).<http://doi.org/10.1109/CSIE.2009.592>
- [21] Plagge, M. (2013). Using Artificial Neural Networks to Predict First-year Traditional Students Second Year Retention Rates. In *Proceedings of the 51st ACM Southeast Conference* (pp. 17:1–17:5). New York, NY, USA: ACM.<http://doi.org/10.1145/2498328.2500061>
- [22] Vihavainen, A. (2013). Predicting Students' Performance in an Introductory Programming Course Using Data from Students' Own Programming Process. In *Advanced Learning Technologies (ICALT), 2013 IEEE 13th International Conference on* (pp. 498–499).<http://doi.org/10.1109/ICALT.2013.161>
- [23] Wolff, A., Zdrahal, Z., Nikolov, A., and Pantucek, M. (2013). Improving Retention: Predicting At-risk Students by Analysing Clicking Behaviour in a Virtual Learning Environment. In *Proceedings of the Third International Conference on Learning Analytics and Knowledge* (pp. 145–149). New York, NY, USA: ACM.<http://doi.org/10.1145/2460296.2460324>
- [24] Azmi, M. S.B.M., and Paris, I. H. B. M. (2011). Academic performance prediction based on voting technique. In *Communication Software and Networks (ICCSN), 2011 IEEE 3rd International Conference on* (pp. 24–27).<http://doi.org/10.1109/ICCSN.2011.6014841>
- [25] Jiang, S., Williams, A., Schenke, K., Warschauer, M., and O'dowd, D. (2014). Predicting MOOC performance with week 1 behavior. In *Educational Data Mining 2014*.
- [26] Ashenafi, M. M., Ronchetti, M., and Riccardi, G. (2016). Predicting student progress from peer-assessment data. In *The 9th International Conference on Educational Data Mining (EDM), 2016 IEDMS*
- [27] Abdel-Salam, T., Kauffmann, P., and Williamson, K. (2005). A case study: do high school GPA/SAT scores predict the performance of freshmen engineering students? In *Frontiers in Education, 2005. FIE '05. Proceedings 35th Annual Conference* (p. S2E–7).<http://doi.org/10.1109/FIE.2005.1612225>
- [28] Golding, P., and McNamarah, S. (2005). Predicting Academic Performance in the School of Computing & Information Technology (SCIT). In *Frontiers in Education, 2005. FIE '05. Proceedings 35th Annual Conference* (p. S2H–S2H).<http://doi.org/10.1109/FIE.2005.1612248>
- [29] Trstenjak, B., and Donko, D. (2014). Determining the impact of demographic features in predicting student success in Croatia. In *Information and Communication Technology, Electronics and Microelectronics (MIPRO), 2014 37th International Convention on* (pp. 1222–1227).<http://doi.org/10.1109/MIPRO.2014.6859754>
- [30] Mishra, T., Kumar, D., and Gupta, S. (2014). Mining Students' Data for Prediction Performance. In *Advanced Computing Communication Technologies (ACCT), 2014 Fourth International Conference on* (pp. 255–262).<http://doi.org/10.1109/ACCT.2014.105>
- [31] Sarker, F., Tiropanis, T., and Davis, H. C. (2014). Linked data, data mining and external open data for better prediction of at-risk students. In *Control, Decision and Information Technologies (CoDIT), 2014 International Conference on* (pp. 652–657).<http://doi.org/10.1109/CoDIT.2014.6996973>
- [32] Christian, T. M., and Ayub, M. (2014). Exploration of classification using NBTree for predicting students' performance. In *Data and Software Engineering (ICODSE), 2014 International Conference on* (pp. 1–6).<http://doi.org/10.1109/ICODSE.2014.7062654>
- [33] Gulati, H. (2015). Predictive analytics using data mining technique. In *Computing for Sustainable Global Development (INDIACom), 2015 2nd International Conference on* (pp. 713–716).

- [34] Fernandez-Delgado, M., Mucientes, M., Vazquez-Barreiros, B., and Lama, M. (2014). Learning analytics for the prediction of the educational objectives achievement. In *Frontiers in Education Conference (FIE), 2014 IEEE* (pp. 1–4).<http://doi.org/10.1109/FIE.2014.7044402>
- [35] Amnueypornsakul, B., Bhat, S., and Chinpruthiwong, P. (2014). Predicting Attrition along the Way: The UIUC Model. *EMNLP 2014*, 55.
- [36] Kloft, M., Stiehler, F., Zheng, Z., and Pinkwart, N. (2014). Predicting MOOC dropout over weeks using machine learning methods. *EMNLP 2014*, 60.
- [37] Bydovska, H., and Popelinsky, L. (2013). Predicting Student Performance in Higher Education. In *Database and Expert Systems Applications (DEXA), 2013 24th International Workshop on* (pp. 141–145).<http://doi.org/10.1109/DEXA.2013.22>
- [38] Minaei-Bidgoli, B., Kashy, D. A., Kortemeyer, G., and Punch, W. F. (2003). Predicting student performance: an application of data mining methods with an educational Web-based system. In *Frontiers in Education, 2003. FIE 2003 33rd Annual (Vol. 1, p. T2A–13)*.<http://doi.org/10.1109/FIE.2003.1263284>
- [39] Slim, A., Heileman, G. L., Kozlick, J., and Abdallah, C. T. (2014). Employing Markov Networks on Curriculum Graphs to Predict Student Performance. In *Machine Learning and Applications (ICMLA), 2014 13th International Conference on* (pp. 415–418).<http://doi.org/10.1109/ICMLA.2014.74>
- [40] Elbadrawy, A., Studham, R. S., and Karypis, G. (2015). Collaborative Multi-regression Models for Predicting Students' Performance in Course Activities. In *Proceedings of the Fifth International Conference on Learning Analytics and Knowledge* (pp. 103–107). New York, NY, USA: ACM.<http://doi.org/10.1145/2723576.2723590>
- [41] Wang, R., Harari, G., Hao, P., Zhou, X., and Campbell, A. T. (2015). SmartGPA: How Smartphones Can Assess and Predict Academic Performance of College Students. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (pp. 295–306). New York, NY, USA: ACM.<http://doi.org/10.1145/2750858.2804251>
- [42] Chaplot, D. S., Rhim, E., and Kim, J. (2015). Predicting student attrition in moocs using sentiment analysis and neural networks. In *Proceedings of AIED 2015 Fourth Workshop on Intelligent Support for Learning in Groups*.
- [43] Watson, C., Li, F. W. B., and Godwin, J. L. (2013). Predicting Performance in an Introductory Programming Course by Logging and Analyzing Student Programming Behavior. In *Advanced Learning Technologies (ICALT), 2013 IEEE 13th International Conference on* (pp. 319–323).<http://doi.org/10.1109/ICALT.2013.99>
- [44] Meier, Y., Xu, J., Atan, O., and van der Schaar, M. (2015). Predicting Grades. *Signal Processing, IEEE Transactions on*, PP (99), 1.<http://doi.org/10.1109/TSP.2015.2496278>
- [45] Jordan, K. (2013). MOOC Completion Rates: The Data, Available at: <http://www.katyjordan.com/MOOCproject.html> [Accessed: 01/08/16]
- [46] Yang, D., Sinha, T., Adamson, D., and Rose, C. P. (2013). Turn on, tune in, and drop out: Anticipating student dropouts in massive open online courses. In *Proceedings of the 2013 NIPS Data-Driven Education Workshop (Vol. 11, p. 14)*.
- [47] Sharkey, M., and Sanders, R. (2014). A Process for Predicting MOOC Attrition. *EMNLP 2014*, 50.
- [47] Ye, C., and Biswas, G. (2014). Early Prediction of Student Dropout and Performance in MOOCs using Higher Granularity Temporal Information. *Journal of Learning Analytics*, 1(3), 169-172.

- [48] Sinha, T., Li, N., Jermann, P., and Dillenbourg, P. (2014). Capturing " attrition intensifying" structural traits from didactic interaction sequences of MOOC learners. EMNLP 2014, 42.
- [49] Coleman, C. A., Seaton, D. T., and Chuang, I. (2015). Probabilistic Use Cases: Discovering Behavioral Patterns for Predicting Certification. In Proceedings of the Second (2015) ACM Conference on Learning @ Scale (pp. 141–148). New York, NY, USA: ACM.<http://doi.org/10.1145/2724660.2724662>
- [50] Kennedy, G., Coffrin, C., de Barba, P., and Corrin, L. (2015). Predicting Success: How Learners' Prior Knowledge, Skills and Activities Predict MOOC Performance. In Proceedings of the Fifth International Conference on Learning Analytics and Knowledge (pp. 136–140). New York, NY, USA: ACM.<http://doi.org/10.1145/2723576.2723593>
- [51] Gedeon, T. D., and Turner, S. (1993). Explaining student grades predicted by a neural network. In Neural Networks, 1993. IJCNN '93-Nagoya. Proceedings of 1993 International Joint Conference on (Vol. 1, pp. 609–612 vol.1).<http://doi.org/10.1109/IJCNN.1993.713989>
- [52] Kotsiantis, S. B., and Pintelas, P. E. (2005). Predicting students marks in Hellenic Open University. In Advanced Learning Technologies, 2005. ICALT 2005. Fifth IEEE International Conference on (pp. 664–668).<http://doi.org/10.1109/ICALT.2005.223>
- [53] Golding, P., and Donaldson, O. (2006). Predicting Academic Performance. In Frontiers in Education Conference, 36th Annual (pp. 21–26).<http://doi.org/10.1109/FIE.2006.32266>
- [54] Ogor, E. N. (2007). Student Academic Performance Monitoring and Evaluation Using Data Mining Techniques. In Electronics, Robotics and Automotive Mechanics Conference, 2007. CERMA 2007 (pp. 354–359).<http://doi.org/10.1109/CERMA.2007.4367712>
- [55] Alenezi, J. K., Awany, M. M., and Fahmy, M. M. M. (2009). Effectiveness of Artificial Neural Networks in forecasting failure risk for pre-medical students. In Computer Engineering Systems, 2009. ICCES 2009. International Conference on (pp. 135–138).<http://doi.org/10.1109/ICCES.2009.5383294>
- [56] García, E. P. I., and Mora, P. M. (2011). Model Prediction of Academic Performance for First Year Students. In Artificial Intelligence (MICAI), 2011 10th Mexican International Conference on (pp. 169–174).<http://doi.org/10.1109/MICAI.2011.28>
- [57] Schalk, P. D., Wick, D. P., Turner, P. R., and Ramsdell, M. W. (2011). Predictive assessment of student performance for early strategic guidance. In Frontiers in Education Conference (FIE), 2011 (pp. S2H-1–S2H-5).<http://doi.org/10.1109/FIE.2011.6143086>
- [58] Gamulin, J., Gamulin, O., and Kermek, D. (2013). Data mining in hybrid learning: Possibility to predict the final exam result. In Information Communication Technology Electronics Microelectronics (MIPRO), 2013 36th International Convention on (pp. 591–596)
- [59] Zheng, J., Chen, Z., and Zhou, C. (2013). Applying NN-based data mining to learning performance assessment. In Conference Anthology, IEEE (pp. 1–5).<http://doi.org/10.1109/ANTHOLOGY.2013.6784924>
- [60] Sharabiani, A., Karim, F., Sharabiani, A., Atanasov, M., and Darabi, H. (2014). An enhanced bayesian network model for prediction of students' academic performance in engineering programs. In Global Engineering Education Conference (EDUCON), 2014 IEEE (pp. 832–837).<http://doi.org/10.1109/EDUCON.2014.6826192>
- [61] Sorour, S. E., Mine, T., Godaz, K., and Hirokawax, S. (2014). Comments Data Mining for Evaluating Student's Performance. In Advanced Applied Informatics (IIAIAI), 2014 IIAI 3rd International Conference on (pp. 25–30).<http://doi.org/10.1109/IIAI-AAI.2014.17>