Improving Early Dropout Detection in Undergraduate Students: Exploring Key Predictors through SHAP Values

Winn Chow
The University of Melbourne
Corresponding Author Email: winn.chow1@unimelb.edu.au

ABSTRACT

CONTEXT

Despite substantial government investments in higher education, students encounter considerable hurdles in completing their educational journeys. In Australia, a concerning 25.4% of students who started their degree studies in 2017 had dropped out by the end of 2022. This problem of high dropout rates extends globally. However, research on the implementation and effectiveness of early warning systems in higher education remains insufficient.

PURPOSE OR GOAL

Ideally, an early warning system should pinpoint students at risk of dropping out as early as their enrolment. However, there is a noticeable gap in research regarding how the indicators of at-risk students evolve over time from enrolment onwards. This study aimed to investigate the use of machine learning, alongside explainable AI, to identify and examine key predictors for the early detection of student dropout at different stages of their first year of study.

APPROACH OR METHODOLOGY/METHODS

We analysed student data from various undergraduate courses at the Polytechnic Institute of Portalegre, Portugal, spanning the academic years from 2008-09 to 2018-19. This dataset included 17 relevant features. The availability of these features varied depending on the stage: (1) admission, (2) the end of 1st semester, and (3) the end of 2nd semester. At these stages of students' first-year study, we developed random forest models to predict dropout cases. Notably, we explored the use of explainable AI, specifically SHAP values, to gain insights into the key features of these models and their evolution across the three stages.

ACTUAL OR ANTICIPATED OUTCOMES

Throughout the three stages, the performance of the random forest models consistently improved. Initially, during the admission stage, the prediction performance was relatively poor (F1 = 0.68). However, optimal performance (F1 = 0.85) was achieved when incorporating the 2nd semester academic results in the prediction. Additionally, there was a notable shift in predictive factors for dropout. Features that held significant importance at the admission stage gradually diminished in significance, while the most recent academic performance features at each stage always surpassed other features in importance.

CONCLUSIONS/RECOMMENDATIONS/SUMMARY

This study investigated the evolution of key predictors of student dropout over time. It was the first to use SHAP values to analyse the shifting importance of individual predictors. These findings enhance our understanding of the dynamic nature of predictors for student dropout risk and aid in applying targeted early intervention strategies for at-risk students.

KEYWORDS

Early dropout detection, Key predictors, Machine learning, Explainable AI, SHAP values

Introduction

Higher education aims to empower individuals with the knowledge, skills, and values necessary to lead fulfilling lives, contribute meaningfully to society, and address the complex challenges facing humanity in the 21st century. Despite substantial government investments in higher education, students encounter considerable hurdles in completing their educational journeys. For instance, the Australian government allocated a record level of annual funding for higher education, amounting to $20 billion in 2022-23 (Department of Education, 2023). However, it is alarming that Australian university degree dropouts also reached a record high in 2022. According to the Australian Federal Education Department (Hare, 2024), 25.4 percent of students who began their studies in 2017 had dropped out by the end of 2022. This dropout rate marks the highest recorded since 2005, with a 1.3 percentage point increase from the previous corresponding period. The issue of high dropout rates is indeed not unique to any particular country but is rather a global phenomenon. The Organization for Economic Co-operation and Development (OECD) reports an average university dropout rate of 32% for the 2020 cohort across 38 OECD countries from North and South America to Europe and Asia-Pacific (OECD, 2022).

Although high dropout rates do not necessarily indicate an inadequate tertiary system, as students may leave courses for diverse reasons such as securing enticing employment opportunities before completion, those who unfortunately fail to complete their courses due to poor academic performance endure a significant personal cost. In addition to the high financial burden of an incomplete course, dropout students earn lower salaries due to the lack of education, endure wasted time, and experience psychological feelings of failure (Hare, 2024).

In recent years, many countries have implemented a range of policies aimed at increasing tertiary completion rates (OECD, 2022). One common approach is to partially tie institutional funding to student completion rates. Other policies concentrate on assisting students in making informed decisions about their field of study.

Among the various prevention initiatives aimed at supporting students to complete their studies at higher education institutions, early warning systems for the early identification of academically at-risk students are recognized as holding significant potential. Early warning systems can be defined as systems “based on student data to identify students who exhibit behaviour or academic performance that puts them at risk of dropping out of school” (U.S. Department of Education, 2016). Such a system can help higher education institutions identify individual at-risk students early and offer targeted support and early interventions.

In recent years, there has been a gradual adoption of early warning systems in higher education, particularly in the United States (Plak et al., 2022). However, research on the use and evaluation of early warning systems in higher education remains scarce and inadequate. In particular, the identification of key predictors for early identification of academically at-risk students in higher education has produced insufficient findings and remains inconclusive (Lee & Chung, 2019; Rowtho, 2017).

Ideally, an early warning system should identify students likely to drop out as early as they are enrolled in the institution. It should then continue to monitor the academic performance of these students, especially during their first year of study, and regularly reassess their chance of dropping out. However, there is a notable gap in research regarding how the key predictors of at-risk students evolve as more information is gathered about them through assessments from enrolment to the conclusion of their first year of study. Although some initial results have been reported (Aulck et al., 2019; Kiss et al., 2019), a comprehensive understanding of how the key predictors of at-risk students evolve remains elusive. This research gap hampers our understanding of the dynamic nature of student dropout risk predictors and constrains the improvement of early identification and intervention for at-risk students during the early stages of their academic journey.

This study investigated the use of machine learning and explainable AI to detect key predictors for early identification of student dropout and how these key predictors evolved in importance during their first year of undergraduate study. We analysed student data from various undergraduate courses at the Polytechnic Institute of Portalegre, Portugal, spanning the academic years from 2008-09 to 2018-19 (Martins et al., 2021). We constructed several prediction models using decision tree-based machine learning techniques at three distinct stages of students' first-year study, from enrolment to the conclusion of their first year, to predict dropout cases. These models were subsequently assessed and compared based on their prediction precision, recall, and F1-score. We further explored the use of explainable AI, specifically Shapley values, to detect the key features of these models.

Here is the structure of the paper. Firstly, a review of relevant literature is presented. Secondly, the data and methodology are introduced. Thirdly, the prediction results from different prediction models are analysed. Fourthly, the results regarding the key predictors of student dropout and how they changed in significance are discussed. Lastly, the paper concludes with key findings and implications.

Literature Review

Machine learning remains the state-of-the-art approach for making predictions based on diverse datasets. The prediction of student dropout is often approached as a classification problem in machine learning, aiming to classify whether a student is at risk of dropping out or not (Martins et al., 2021; Nagy & Molontay, 2018). A wide range of machine learning techniques has been applied to predict student dropout, including logistic regression (Nagy & Molontay, 2018), support vector machine (Pallathadka et al., 2023), and decision trees (Lee & Chung, 2019; Martins et al., 2021). In comparison with the different machine learning techniques, decision tree-based techniques, specifically random forest and boosted trees, are often shown to achieve the best prediction accuracy.

Random forest and boosted trees, such as Extreme Gradient Boosting (XGBoost), are frequently the preferred choices for tackling classification problems due to their distinctive characteristics. These methods generally achieve high prediction accuracy by employing ensemble learning techniques, which reduce overfitting by combining predictions from multiple individual models.

The selection of data features for predicting student dropout is crucial for the performance of the prediction models. When examining data features available when students enrol in a university, certain features are often observed to significantly enhance predictive capability. Among demographic variables, age at enrolment and gender are particularly important features. Specifically, Chen et al. (2018) observed that the enrolled ages of students influenced dropout prediction, noting that older students were more likely to graduate in the STEM fields. In Rowtho's study (2017) on undergraduate student performance prediction, gender emerged as a significant predictor of GPA.

Mixed results arise when assessing whether students' domestic or international status significantly predicts the likelihood of dropout, although relevant studies are limited. After investigating the at-risk program at a university, Dobele et al. (2013) discovered that domestic students were significantly more likely to be at risk of dropout than international students. However, according to Education Inspectorate (Upton, 2022), the dropout rate was about four times higher among international students than Dutch students in the Netherlands.

In terms of socio-economic factors, student debt or financial status is identified as a significant predictor of dropout in several studies. For instance, Bello et al. (2020) discovered that low family income significantly increased the likelihood of dropout among first-year Informatics Engineering students. Similarly, Thammasiri et al. (2014) identified student loan among the top 10 important features for predicting freshman student dropout.

Considering the pre-entry academic performance, the admission grade or high school GPA emerges as a significant predictor of dropout, especially in cases where post-entry university academic performance is not considered. In a study focusing on predicting university dropout based solely on pre-entry information, Nagy and Molontay (2018) consistently identified admission points as one of the most crucial features across four prediction models.

After students enrol in a university, numerous studies consistently indicate that university academic performance significantly predicts student dropout. For example, Chen et al. (2018) found that GPA in the first and second semesters strongly predicted dropout across STEM majors. This connection is intuitive, as university academic performance directly reflects students' learning abilities and their capacity to meet course requirements, with poor academic performance typically being the primary cause of university dropout.

Other potential predictors, such as being a scholarship holder or having special education needs, have been examined only in a few studies. In Thammasiri et al.'s investigation of freshman student dropout (2014), tuition waiver scholarship holder emerged as a significant factor for predicting dropout. Additionally, Rußmann et al. (2023) found that students with mental health, learning, and physical disabilities were significantly more inclined to consider dropping out of higher education based on a 2020 Germany-wide student survey.

Research examining how key predictors of student dropout evolve from enrolment to the end of their first year of study, as more assessment data about students becomes available, is lacking. Recent findings, such as those from Kiss et al.'s investigation focusing on STEM students (2019) and Aulck et al.'s analysis of first-year undergraduate students at a US university (2019), indicate that prediction models relying solely on demographic data and pre-entry information demonstrate relatively weaker predictive abilities for dropout. Notably, Kiss et al.'s study (2019) even incorporating only the first week of student performance at university into the model could moderately enhance dropout prediction. Furthermore, both studies achieved significantly higher prediction accuracy when they included the students' first-year academic performance in their prediction models. Kiss et al. (2019) identified certain pre-entry information, such as age and admission score, as remaining important predictors even when post-entry academic performance was considered in the prediction model. In contrast, Aulck et al. (2019) found that a prediction model based solely on a summary of first-year performance data yielded the best predictive performance.

Understanding the role and importance of key predictors in an advanced machine learning-based prediction model, such as random forest and XGBoost, is crucial for gaining insights into the prediction process. This understanding facilitates the optimization of model performance. However, these advanced machine learning techniques, which boast higher performance, are often built on complex algorithms and thus face criticism for their lack of explainability. For example, random forest and XGBoost operate as ensemble methods, combining multiple decision trees to make predictions, which makes it challenging to interpret the exact reasoning behind individual predictions. Nevertheless, recent research and development efforts have improved the explainability of these models, particularly in the realm of explainable AI, with techniques such as feature importance rankings, partial dependence plots, and Shapley values (Saranya & Subhashini, 2023). Despite these advancements, few studies have applied these techniques to investigate the contribution and importance of key predictors of student dropout, possibly owing to their recency. For instance, Beaulac and Rosenthal (2019) and Kiss et al. (2019) solely applied permutation-based feature importance ranking to analyse their random forest and XGBoost prediction models in their study of student dropouts. However, feature importance ranking—an earlier development—only provides a ranking of features based on their importance in the model, without indicating whether a feature is positively or negatively contributing to the prediction.

Shapley values offer a more sophisticated and comprehensive approach to understanding feature importance in machine learning models. Originating from cooperative game theory, they offer a strong theoretical foundation and have subsequently been adapted for use in interpreting machine learning models (Munn & Pitman, 2022). In essence, Shapley values quantify the contribution of a feature to a model's specific prediction by evaluating its marginal contribution across all possible combinations with other features. A positive Shapley value for a feature indicates that its presence increases the model's prediction compared to the average prediction, while a negative Shapley value indicates the opposite. Additionally, Shapley values can provide a global interpretation of feature importance by assessing the impact of a feature across the entire dataset. However, to our best knowledge, research on student dropout has yet to adopt Shapley values for evaluating the contributions of individual key predictors.

Data and Methodology

The institutional data concerning undergraduate students enrolled at the Polytechnic Institute of Portalegre, Portugal was analysed in this study. The dataset used in this analysis was publicly accessible on the UC Irvine machine learning repository (Realinho et al., 2021). It encompassed student records from the academic years 2008/09 to 2018/2019, across various undergraduate programs including social service, management, and informatics engineering. The dataset contained 4,424 student records, all of which were complete with no missing values. A total of 17 relevant features were utilized for analysis, categorized into demographic, socio-economic, pre-entry admission, post-entry university academic performance, and final dropout status. This dataset was selected because it tracks students from admission through their first year, enabling an analysis of how at-risk indicators evolve over time.

Demographic-related features included gender, age at enrolment, international status, and special education needs. Socio-economic-related features encompassed debtor status and up-to-date tuition fees. Pre-entry admission-related information comprised application order, attendance mode, scholarship holder status, and admission grade. Post-entry university academic performance-related features included enrolled and approved curricular units for both the first and second semesters, as well as the average grades. The final feature indicated whether the student dropped out at the conclusion of the course.

Firstly, we examined the data. We observed a common occurrence often seen in student dropout datasets: an imbalance in the dropout status. Specifically, approximately 32% of the students were classified as dropouts, while the remaining 68% were categorized as non-dropouts. This imbalance posed a problem for dropout classification that required handling.

To address this data imbalance, we adopted a commonly recommended approach: over-sampling with SMOTE (Chawla et al., 2002). This technique, extensively used in various dropout classification studies (Lee & Chung, 2019; Martins et al., 2021), helped mitigate the imbalance by generating synthetic samples for the minority class. By doing so, it enhanced the models' capability to learn from and accurately classify instances of the minority class. Specifically, we applied SMOTE as a preprocessing step on the training data before training the prediction models. Notably, SMOTE was not applied to the test data to ensure an accurate assessment of testing performance.

We then proceeded to construct prediction models to evaluate their effectiveness in predicting dropouts at three distinct stages of students' early academic journey: (1) admission, (2) the end of the 1st semester of the first year of study, and (3) the end of the 2nd semester of the first year of study. At each stage, we developed one prediction model using random forest, chosen for its consistently superior performance in dropout prediction studies (Lee & Chung, 2019; Martins et al., 2021). Additionally, we tailored the sets of features used for prediction based on their availability at each stage. For instance, only pre-entry admission features were used to train the prediction model at the admission stage.

To prepare the data for training, we followed the common practice of dividing the dataset into a training set (80%) and a test set (20%). As previously mentioned, we applied SMOTE on the training set to ensure the balance of the training data. The training set served both for training and validation purposes. We trained the random forest models using this training set and then evaluated them using the test set. For each prediction model, we utilized a 5-fold cross-validation approach for hyperparameter optimization. We employed grid search cross-validation to explore and identify the optimal set of hyperparameters for the prediction models. To evaluate the prediction performance of the models, we assessed precision, recall, and F1 score, which captured the balance between precision and recall.

Table 1: The dropout prediction results of the test set

Stage	Prediction model	Precision	Recall	F1 score
Admission	Random Forest	0.55	0.70	0.68
End of 1st Semester	Random Forest	0.75	0.73	0.80
End of 2nd Semester	Random Forest	0.80	0.81	0.85

Table 1 displays the dropout prediction results of the test set. Initially, the prediction performance was relatively poor during the admission stage, achieving an F1 score of 0.68. However, a significant enhancement in prediction performance was observed upon incorporating the 1st semester academic results into the prediction model, leading to an F1 score improvement to 0.80, which reflected approximately an 18% increase. The optimal performance was achieved when the 2nd semester academic results and up-to-date tuition fees status were also included in the prediction. This resulted in a further improvement in the F1 score to 0.85, representing approximately a 6% increase from the end of the 1st semester stage.

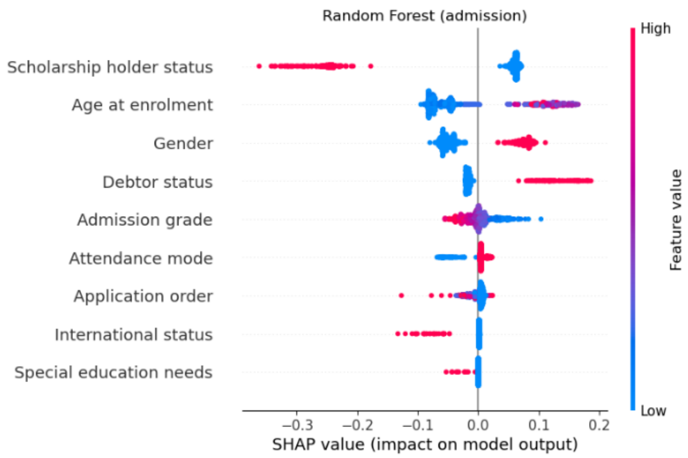

(a) admission

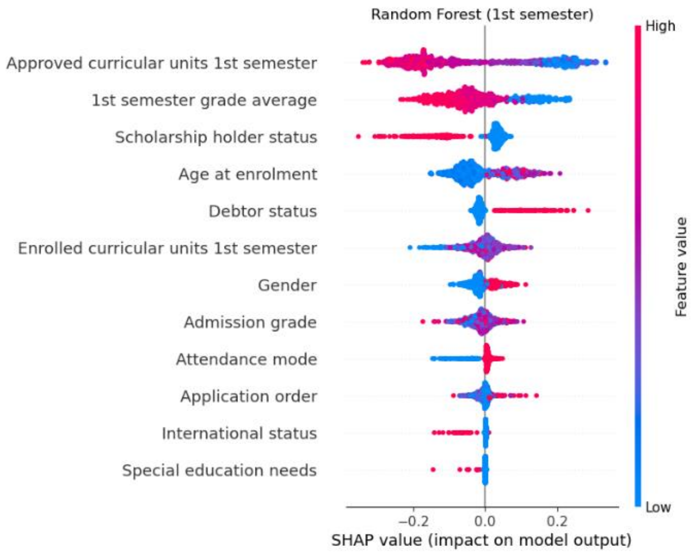

(b) the end of 1st semester

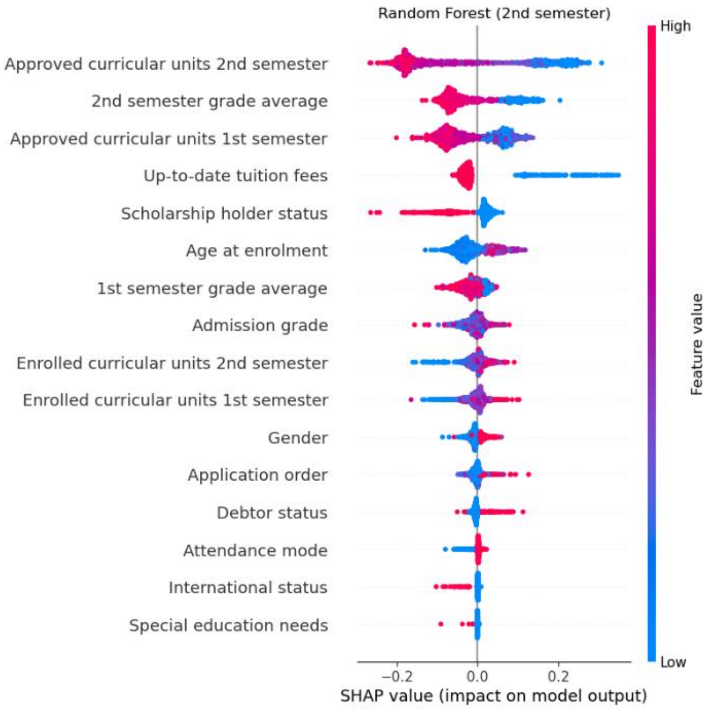

(c) the end of 2nd semester

Figure 1: SHAP values of all features of the random forest prediction models at each stage

Lastly, and perhaps most importantly, we identified the most significant features of the prediction models at various stages and evaluated how their importance changed throughout these stages. We calculated Shapley values for all features across all stages of the prediction models using the test set. The SHAP framework, also referred to as SHAP values (Lundberg & Lee, 2017), facilitated this computation. In Figure 1, you can examine the distribution of SHAP values (displayed on the horizontal axis) calculated across the test set for all features at each stage. Within each stage, the features were arranged based on their importance, with the most influential ones appearing at the top of the plot. Additionally, it illustrates the general relationship between the features and predictions. Features exhibiting positive SHAP values when their feature values were high (depicted in red) and negative SHAP values when their feature values were low (depicted in blue) showcased a positive association with dropout prediction. This suggests that higher values in these features elevated the likelihood of dropout. Conversely, features displaying negative SHAP values when their feature values were high and positive SHAP values when their feature values were low demonstrated a negative association with dropout prediction.

Figure 2 illustrates how the importance of different features evolved as students progressed from admission to the end of the 2nd semester, with increasing availability of academic performance data. This importance was measured by calculating the average absolute SHAP value across the test set, with higher values indicating greater importance. Notably, features such as scholarship holder status, age at enrolment, gender, and debtor status, which held significant importance at the admission stage, gradually diminished in significance as students progressed towards the end of the 2nd semester. By the end of the 1st semester, two academic performance features emerged as the most important: the number of approved curricular units in the 1st semester and the 1st semester grade average. These features surpassed all other features in importance at admission, indicating a notable shift in predictive factors for dropout once academic performance data became accessible. With more recent academic performance data available by the end of the 2nd semester, the most important features shifted once more, with the emergence of two 2nd semester academic performance metrics as the most influential: the number of approved curricular units and the 2nd semester grade average. Additionally, the behaviour of students not settling their tuition fees in full also emerged as a significant feature.

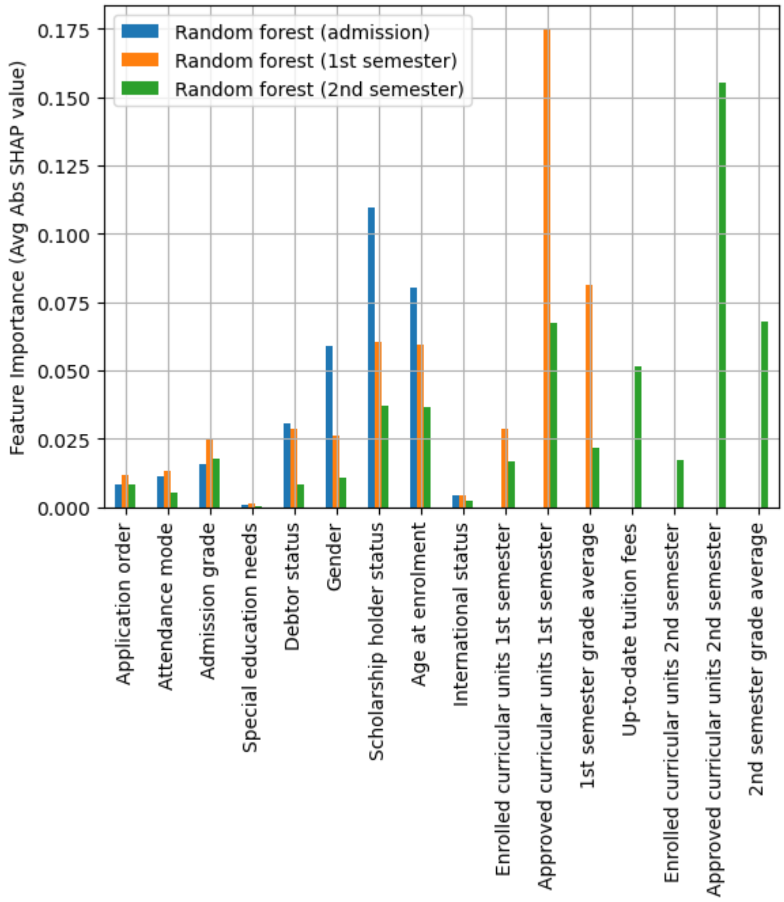

Figure 2: The change in importance of the key features across the three stages

Discussion and Limitations

This study revealed several noteworthy and significant findings. Firstly, the random forest models demonstrated a consistent improvement in performance across the three stages. Initially, during the admission stage, the prediction performance was relatively poor (F1 = 0.68). However, optimal performance (F1 = 0.85) was achieved when incorporating the 2nd semester academic results in the prediction. This result aligns with previous research by Kiss et al. (2019) and Aulck et al. (2019), highlighting that integrating the most recent academic performance can notably enhance dropout prediction. Secondly, the key predictors of dropout at three distinct stages of a student's early academic journey were identified using SHAP values. At the admission stage, scholarship holder status, age at enrolment, and gender emerged as the most important features. Among these features, scholarship holder status was associated with a lower likelihood of dropout, consistent with findings reported by Thammasiri et al. (2014). Conversely, older students and male students were found to be more prone to dropping out across various undergraduate programs. This result contrasts with previous findings that older students were more inclined to graduate in STEM fields (Chen et al., 2018). However, these students might encounter distinct challenges concerning academic preparedness and balancing family obligations or work commitments. At the end of both the 1st and 2nd semesters, the two features that predominantly represented students' most recent academic performance assumed prominence and became the top two most important features. These features were the number of approved curricular units and the grade average. Lastly, the importance of different features evolved as students progressed from admission to the end of the 2nd semester. Notably, features that held significant importance at the admission stage gradually diminished in significance, while the most recent academic performance features at each stage always surpassed other features in importance.

In conclusion, this study investigated the use of machine learning and explainable AI to detect key predictors for early identification of student dropout and how these key predictors evolved in importance during the first year of undergraduate study. This study contributes to our understanding of the dynamic nature of predictors for student dropout risk. To our knowledge, this is the first study to use SHAP values to analyse the shifting importance of individual predictors. SHAP values prove to be instrumental in providing valuable insights into the prediction process, offering a theoretically sound method for quantifying each predictor's contribution to the model's prediction. Practically, these findings improve our comprehension of the evolution in importance of key predictors of student dropouts early in their academic journey, enabling the implementation of targeted early intervention strategies for at-risk students. Given that recent academic performance is the most reliable indicator of dropout risk, continuous monitoring of student performance is recommended to facilitate early identification and intervention. However, it is important to note that these findings are based on one dataset from one university. Further research is needed to validate the results at other universities. Exploring a broader set of dropout risk indicators and studying their evolution over time are also needed. Additionally, we recommend a wider application of SHAP values to assess feature importance in other studies.

References

Aulck, L. S., Nambi, D., Velagapudi, N., Blumenstock, J., & West, J. D. (2019, July 1). *Mining University Registrar Records to Predict First-Year Undergraduate Attrition*. Educational Data Mining. https://www.semanticscholar.org/paper/Mining-University-Registrar-Records-to-Predict-Aulck-Nambi/efb5d83f8281de6bfcca65743bcb185fe13bb6f6

Beaulac, C., & Rosenthal, J. S. (2019). Predicting University Students' Academic Success and Major Using Random Forests. *Research in Higher Education*, *60*(7), 1048–1064. https://doi.org/10.1007/s11162-019-09546-y

Bello, F. A., Kóhler, J., Hinrechsen, K., Araya, V., Hidalgo, L., & Jara, J. L. (2020). Using machine learning methods to identify significant variables for the prediction of first-year Informatics Engineering students dropout. *2020 39th International Conference of the Chilean Computer Science Society (SCCC)*, 1–5. https://doi.org/10.1109/SCCC51225.2020.9281280

Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research*, *16*, 321–357. https://doi.org/10.1613/jair.953

Chen, Y., Johri, A., & Rangwala, H. (2018). Running out of STEM: A comparative study across STEM majors of college students at-risk of dropping out early. *Proceedings of the 8th International Conference on Learning Analytics and Knowledge*, 270–279. https://doi.org/10.1145/3170358.3170410

Department of Education. (2023, November 22). *Australian Government 2022-23 March Budget*. https://www.education.gov.au/about-department/corporate-reporting/budget/202223-march-budget

Dobele, A. R., Gangemi, M., Kopanidis, F., & Thomas, S. (2013). At Risk Policy and Early Intervention Programmes for Underperforming Students: Ensuring Success? *Education & Training*, *55*(1), 69–82. https://doi.org/10.1108/00400911311295022

Hare, J. (2024, January 8). *University degree dropouts reach record*. Australian Financial Review. https://www.afr.com/work-and-careers/education/university-degree-dropouts-reach-record-20240104-p5ev3u

Kiss, B., Nagy, M., Molontay, R., & Csabay, B. (2019). Predicting Dropout Using High School and First-semester Academic Achievement Measures. *2019 17th International Conference on Emerging eLearning Technologies and Applications (ICETA)*, 383–389. https://doi.org/10.1109/ICETA48886.2019.9040158

Lee, S., & Chung, J. Y. (2019). The Machine Learning-Based Dropout Early Warning System for Improving the Performance of Dropout Prediction. *Applied Sciences*, *9*(15), Article 15. https://doi.org/10.3390/app9153093

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 4768–4777.

Martins, M. V., Tolledo, D., Machado, J., Baptista, L. M. T., & Realinho, V. (2021). Early Prediction of student's Performance in Higher Education: A Case Study. In Á. Rocha, H. Adeli, G. Dzemyda, F. Moreira, & A. M. Ramalho Correia (Eds.), *Trends and Applications in Information Systems and Technologies* (pp. 166–175). Springer International Publishing. https://doi.org/10.1007/978-3-030-72657-7_16

Munn, M., & Pitman, D. (2022). *Explainable AI for Practitioners: Designing and Implementing Explainable ML Solutions*. O'Reilly Media, Inc, USA.

Nagy, M., & Molontay, R. (2018). Predicting Dropout in Higher Education Based on Secondary School Performance. *2018 IEEE 22nd International Conference on Intelligent Engineering Systems (INES)*, 000389–000394. https://doi.org/10.1109/INES.2018.8523888

OECD. (2022). *Education at a Glance 2022: OECD Indicators*. Organisation for Economic Co-operation and Development. https://www.oecd-ilibrary.org/education/education-at-a-glance-2022_3197152b-en

Pallathadka, H., Wenda, A., Ramirez-Asís, E., Asís-López, M., Flores-Albornoz, J., & Phasinam, K. (2023). Classification and prediction of student performance data using various machine learning algorithms. *Materials Today: Proceedings*, *80*, 3782–3785. https://doi.org/10.1016/j.matpr.2021.07.382

Plak, S., Cornelisz, I., Meeter, M., & van Klaveren, C. (2022). Early warning systems for more effective student counselling in higher education: Evidence from a Dutch field experiment. *Higher Education Quarterly*, *76*(1), 131–152. https://doi.org/10.1111/hequ.12298

Realinho, V., Vieira Martins, M., Machado, J., & Baptista, L. (2021). *Predict Students' Dropout and Academic Success* [Dataset]. https://doi.org/10.24432/C5MC89

Rowtho, V. (2017). Early Detection of At-Risk Undergraduate Students through Academic Performance Predictors. *Higher Education Studies*, *7*(3), 42–54.

Rußmann, M., Netz, N., & Lörz, M. (2023). Dropout intent of students with disabilities. *Higher Education*. https://doi.org/10.1007/s10734-023-01111-y

Saranya, A., & Subhashini, R. (2023). A systematic review of Explainable Artificial Intelligence models and applications: Recent developments and future trends. *Decision Analytics Journal*, *7*, 100230. https://doi.org/10.1016/j.dajour.2023.100230

Thammasiri, D., Delen, D., Meesad, P., & Kasap, N. (2014). A critical assessment of imbalanced class distribution problem: The case of predicting freshmen student attrition. *Expert Systems with Applications*, *41*(2), 321–330. https://doi.org/10.1016/j.eswa.2013.07.046

Upton, B. (2022, October 23). *Overseas students three times more likely to drop out than Dutch*. Times Higher Education (THE). https://www.timeshighereducation.com/news/overseas-students-three-times-more-likely-drop-out-dutch

U.S. Department of Education. (2016). *Issue Brief: Early Warning Systems*. https://www2.ed.gov/rschstat/eval/high-school/early-warning-systems-brief.pdf

Copyright statement