Machine Learning Approaches to Student Dropout Prediction: A Systematic Review (2020–2024)

Amanda Balnionyte
Vilnius Gediminas Technical University
amanda.balnionyte@gmail.com

VILNIUS TECH
Vilnius Gediminas Technical University

Introduction

Student dropout remains a major issue in higher education, with consequences for both learners and institutions. Early detection of students at risk has therefore become an important goal for universities that want to improve retention. Machine learning supports this by analysing behavioural traces from learning systems, where patterns such as login frequency and interaction with course materials often show early signs of disengagement.
Research in this area has grown. In the screened studies, relevant publications increased from three in 2020 to eight in 2021 and 2022, and reached twelve in 2023 and 2024. This trend shows that dropout prediction has become a more established research direction.
This review examines recent machine learning methods, key features and recurring methodological challenges.

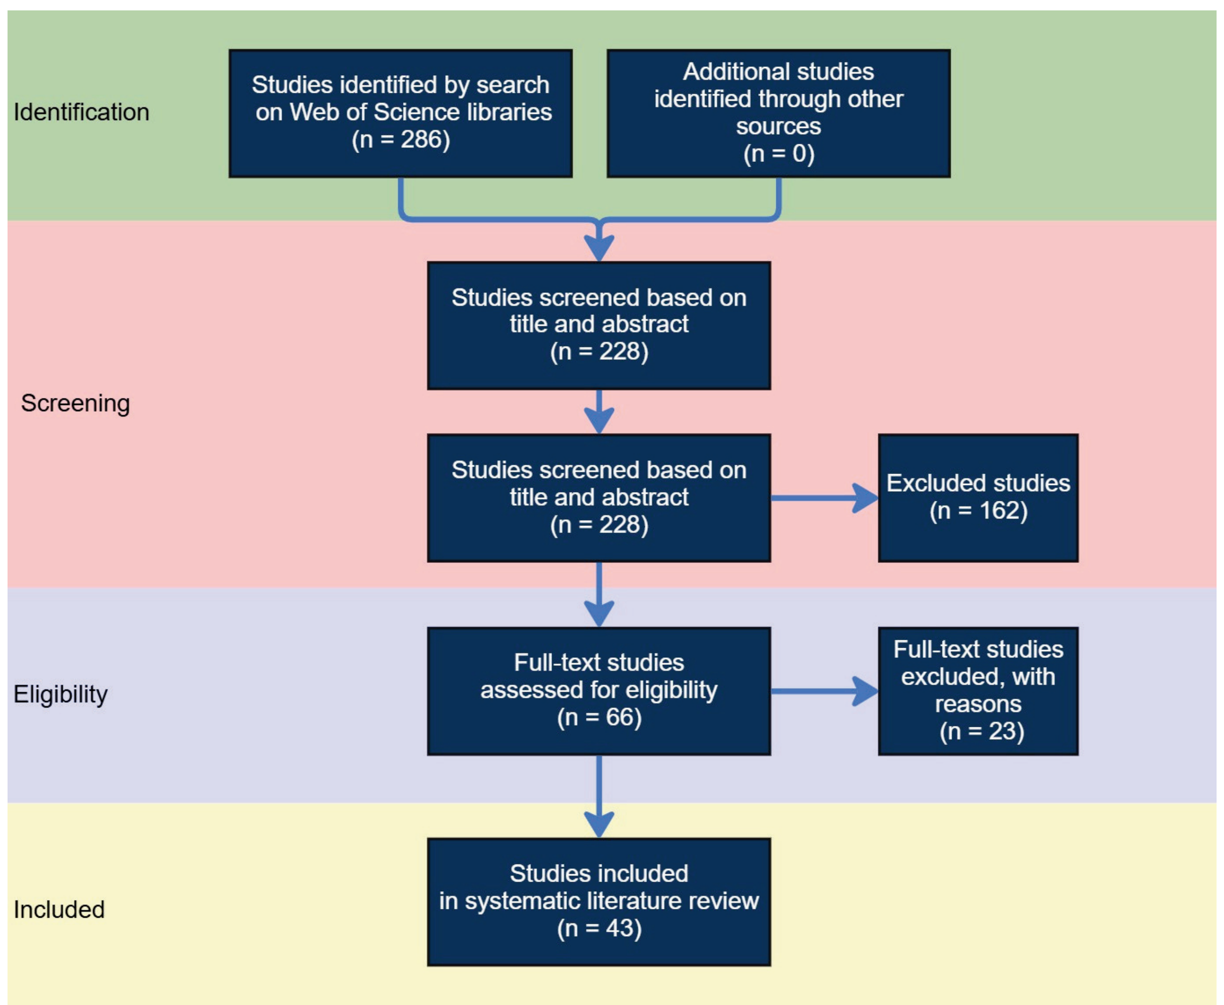

Fig. 1. PRISMA-based study selection process for the systematic review

Methodology

This study uses a Systematic Literature Review to examine how machine learning methods are applied to student dropout prediction in higher education. The review was designed to follow a transparent and replicable process with clear selection criteria. The search was conducted in the Web of Science database, which offers strong coverage of research in computer science, learning analytics and educational technology. Publications from 2020 to 2024 were included to capture recent trends and the growing use of LMS data. Keyword combinations such as “machine learning” AND “dropout”, “learning analytics” AND “higher education”, “LMS logs” AND “prediction” and “Moodle” AND “student dropout” were used to identify relevant studies.
The inclusion criteria required that each study focused on dropout prediction in higher education, used machine learning models and reported empirical evaluation metrics. Studies were excluded when they lacked ML methods, did not provide performance results or focused on non tertiary education. Screening took place in several stages. Duplicates were removed first, followed by title and abstract review. Full texts were then checked to confirm eligibility.
For each selected publication, information was extracted on algorithms, datasets, feature groups and evaluation outcomes. A narrative synthesis approach was used to group studies by methodological similarities and highlight recurring patterns and limitations across the literature.

Results

The structure of the reviewed research becomes clear when examining the selection process shown in Fig. 1. Many publications were removed because they lacked clear evaluation metrics, used incomplete datasets or provided limited detail on feature construction. This filtering ensured that the final analysis was based only on studies with solid methodological foundations. The research landscape shown in Fig. 2 reveals a steady rise in relevant publications from 2020 to 2024, reflecting growing interest in machine learning for early dropout prediction and the increasing availability of LMS data. At the same time, the uneven yearly distribution suggests that the field is still developing and would benefit from more multi institutional evidence.
Algorithm comparison is shown in Fig. 3. Random Forest, Support Vector Machine and Logistic Regression appear most frequently, but accuracy values offer a clearer picture of model performance. The highest accuracy was achieved by Ensemble Methods at 0.984, followed by Polynomial Regression at 0.97, Random Forest at 0.96, Neural Networks at 0.94, Extreme Gradient Boosting at 0.93 and Gradient Boosting Machine at 0.92. The lowest accuracy was 0.75, usually in simpler models or when early semester data contained limited behavioural information.
Feature distribution in Fig. 4 and Fig. 5 highlights that academic performance and learning behaviour dominate across studies. Grades are the most common predictor, followed by LMS activity, assignments, quiz behaviour and page interactions. Together, these features provide a comprehensive view of student performance and support more accurate identification of learners at risk of dropping out.

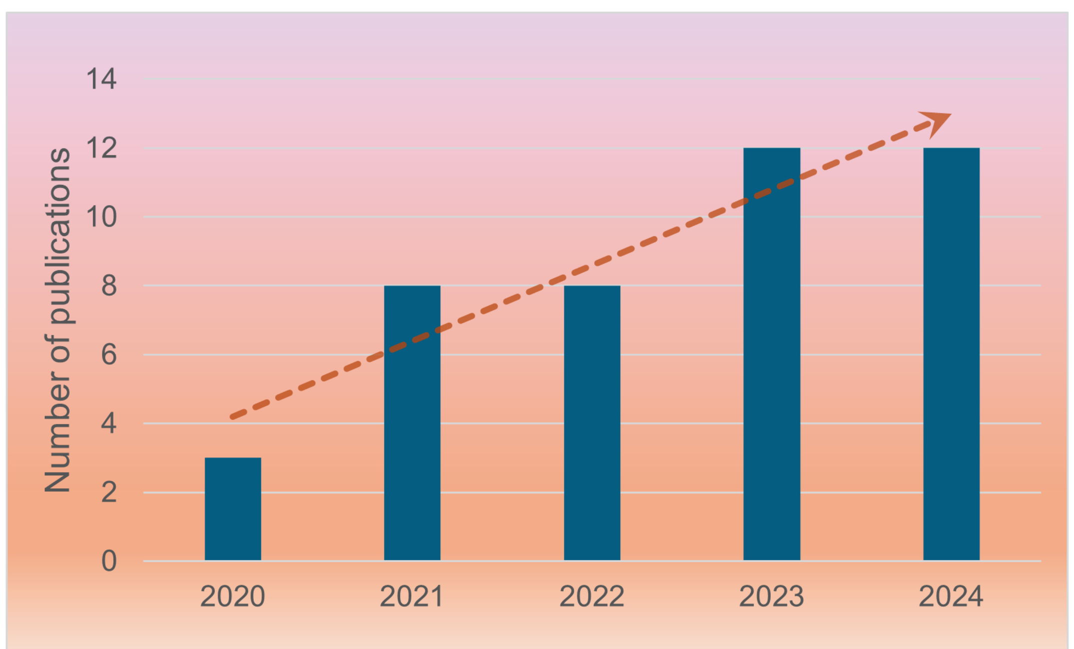

Fig. 2. Final publications per year

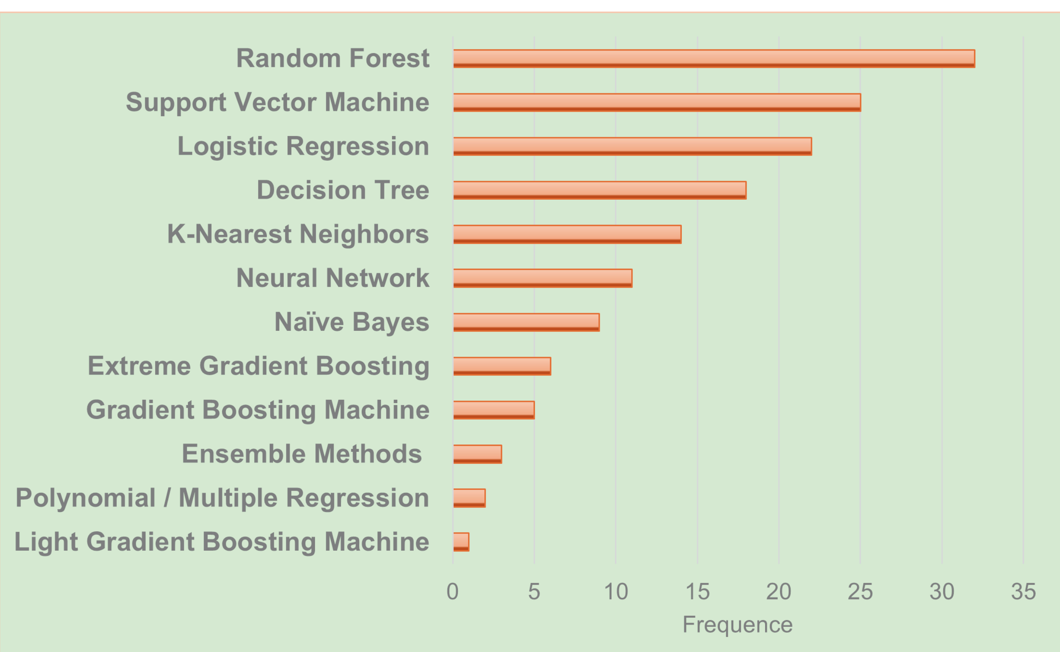

Fig. 3. Frequency of machine learning algorithms

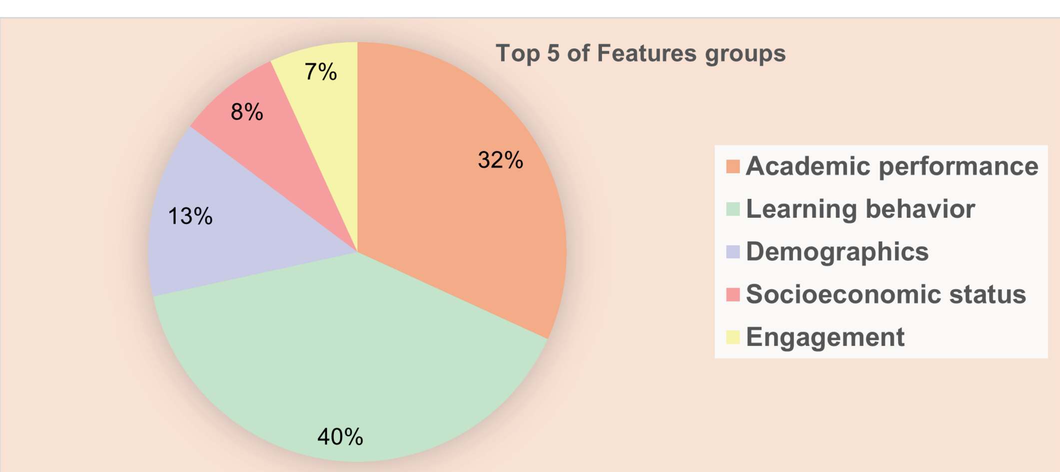

Fig. 4. Distribution of the top five feature groups

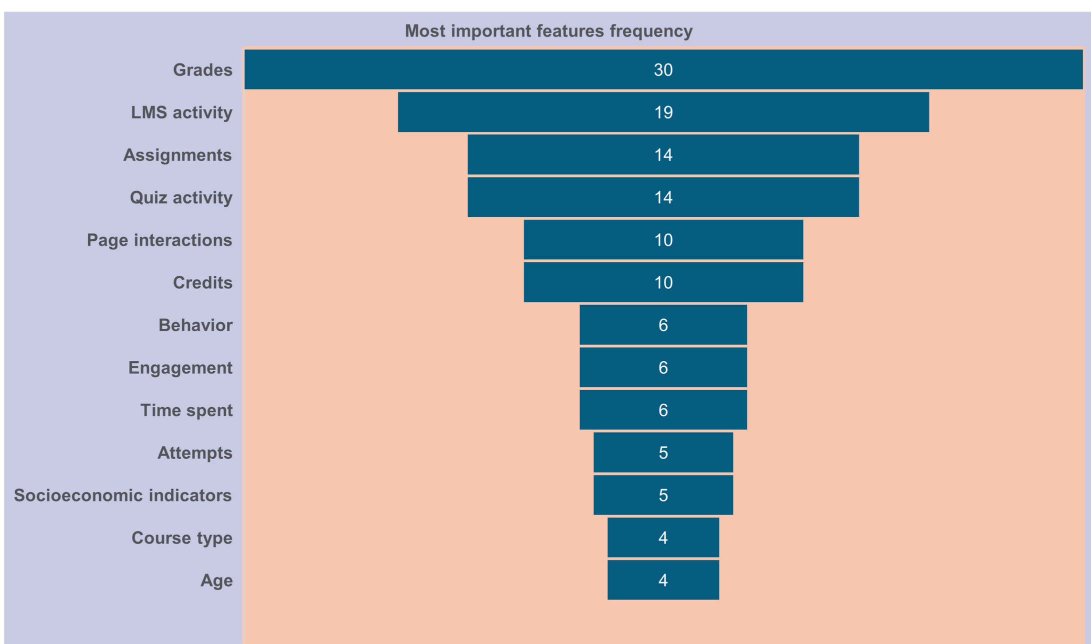

Fig. 5. Frequency of the most important predictive features

Limitations and future directions

Across the reviewed studies, several limitations reduce the strength of the findings. Many publications rely on small or single institution datasets, which limits generalisability and makes it difficult to compare results across different educational settings. Evaluation metrics are also inconsistent, since some studies report only accuracy while others use precision, recall or AUC. Dropout definitions vary between institutions, adding further differences in study design. Only a few papers address transparency, privacy or potential bias, even though these issues are important when predictions influence student support.
Future research should use larger and more diverse datasets, apply consistent evaluation metrics and adopt clearer dropout definitions. Real time LMS data and explainable machine learning methods may improve trust and practical usefulness.

Conclusion

This review shows that machine learning models can provide valuable support for predicting student dropout. Across the analysed studies, Random Forest, Support Vector Machine and Logistic Regression are the most frequently used techniques, yet frequency alone does not reflect performance. Reported accuracy results offer a clearer view of model effectiveness. The highest accuracy was achieved by Ensemble Methods at 0.984, followed by Polynomial Regression at 0.97. Random Forest reached 0.96, Neural Networks achieved 0.94 and Extreme Gradient Boosting and Gradient Boosting Machine followed with 0.93 and 0.92. These results suggest that models able to learn complex relationships tend to perform best.
The strongest predictors remain academic performance indicators such as course grades and accumulated credits, together with behavioural activity recorded in LMS systems, including login patterns and engagement with learning materials. Demographic and financial variables show limited predictive value unless combined with academic or behavioural features. Although results are promising, progress is limited by small datasets, inconsistent metrics and limited attention to transparency. Future studies should address these issues to support more reliable and responsible prediction systems.