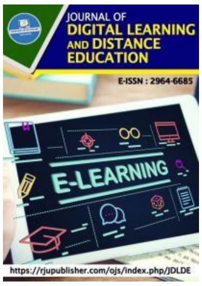

JOURNAL OF DIGITAL LEARNING AND DISTANCE EDUCATION (JDLDE)

ISSN: 2964-6685

https://www.rjupublisher.com/ojs/index.php/JDLDE

Student Success Prediction in Digital Learning Environments

Ejenarhome Otega PROSPER[a], Augustine Osazee AIRHIAVBERE[b], Godfrey Perfectson OISE[b*], Agwam Gladys IFEOMA[c]

[a]*Department of Computer Science, Delta State University, Delta State, Nigeria.*
[b]*Department of Computer Science, University of Benin, Edo State, Nigeria.*
[d]*Department of Computer Science, Western Delta University, Delta State, Nigeria*

*Corresponding author: godfrey.oise@physci.uniben.edu

Abstract

Acknowledging the risk of perpetuating bias in AI-driven student success prediction, this study introduces a fairness-conscious machine learning framework that aims to balance predictive accuracy with ethical responsibility in digital learning settings. Using a dataset of 5,000 anonymized student records, three models, Random Forest (RF), Gradient Boosting (GBM), and Support Vector Machine (SVM) were developed to forecast academic outcomes. Model evaluation combined standard metrics (accuracy, precision, recall, F1-score) with fairness measures such as Demographic Parity, Equal Opportunity, and Disparate Impact Ratio to explore trade-offs between accuracy and fairness. Results indicated that while RF and GBM had slightly higher accuracy, SVM demonstrated more consistent fairness across demographic groups, emphasizing its stronger balance between predictive power and equity. A fairness-centered optimization method was applied to embed fairness constraints directly into model training, showing that both accuracy and fairness can be improved simultaneously rather than being in opposition. The framework integrates fairness throughout data preprocessing, model development, and post-prediction review, promoting transparent and responsible decision-making. By aligning with international ethical AI standards from UNESCO and the OECD, this research provides a practical pathway for creating educational prediction systems that enhance inclusion, minimize bias, and build trust in digital learning environments.

Keywords: Fairness-aware machine learning, Student success prediction, educational data mining, Ethical AI, Digital learning

Acknowledgments: Thank you to all parties who have helped with this research
Ejenarhome Otega PROSPER[a], Augustine Osazee AIRHIAVBERE[b], Godfrey Perfectson OISE[c*], Agwam Gladys IFEOMA[d]

Prosper, E. O., Airhiavbere, A. O., Oise, G. P., & Ifeoma, A. G. (2025). Student Success Prediction in Digital Learning Environments. *Journal of Digital Learning and Distance Education*, *4*(6), 1697-1707. https://doi.org/10.56778/jdlde.v4i6.592

Introduction

Learning Analytics (LA) and Educational Data Mining (EDM) have become essential components of modern higher education as institutions increasingly rely on data-driven approaches to improve teaching, support learners, and enhance retention (Bird et al., 2021). With the growth of digital learning platforms, vast amounts of academic, behavioural, and demographic data are now available for analysis. Machine Learning (ML) and Artificial Intelligence (AI) techniques have been widely adopted to predict student performance, identify at-risk learners, and guide personalized interventions. This shift toward predictive and preventive educational strategies reflects a broader transformation in higher education toward evidence-based decision-making (Bertolini et al., 2021).

While predictive models have demonstrated strong performance across numerous educational datasets, growing attention is being drawn to the ethical, social, and fairness implications associated with AI systems (Bertolini et al., 2022). Predictive algorithms are often trained on historical datasets that reflect existing inequalities associated with socioeconomic background, gender, race, and access to educational resources. Even when sensitive variables are excluded, proxy features may still encode underlying disparities (Uttamchandani & Quick, 2022). As a result, ML models may systematically misclassify certain demographic groups, raising concerns regarding fairness, transparency, and potential reinforcement of educational disadvantage (James et al., 2024). Evidence increasingly shows that ML models performing well at the aggregate level may produce unequal error rates across subgroups, masking disparities in predictive outcomes. Yet, most studies in LA and EDM remain narrowly focused on improving accuracy through hyperparameter tuning, feature engineering, or algorithm comparison (Oise, et al., 2025). Fairness auditing, subgroup evaluation, and bias mitigation remain underdeveloped in the educational domain compared to fields such as healthcare or finance (Oyedotun et al., 2025). This accuracy-centric focus limits the practical and ethical reliability of predictive systems deployed in real educational environments, where fairness is a critical requirement method (Oise, et al., 2025).

Although fairness-aware machine learning methods such as reweighting, fairness constraints, or post-hoc calibration have been successfully applied in other sectors, their adoption in education is limited (Li et al., 2021; Yu et al., 2021).. Few studies provide structured guidance on integrating fairness principles into the entire modelling pipeline, from data pre-processing to model training and post-prediction auditing (Zhao & Ng, 2023). International guidelines, including UNESCO's Recommendation on the Ethics of AI and the OECD AI Principles, emphasize fairness and accountability as essential components of responsible AI deployment. However, the operationalization of these principles in educational predictive modelling remains a significant gap that this study seeks to address,(H. Li et al., 2024; Sha et al., 2023).

To respond to these gaps, this study investigates the integration of fairness-aware machine learning into student success prediction using a dataset of 5,000 de-identified records. The study pursues four key objectives: to develop and evaluate three machine learning models Random Forest, Gradient Boosting, and Support Vector Machine for predicting academic outcomes; to assess each model's fairness across demographic subgroups using Demographic Parity, Equal Opportunity, and Disparate Impact Ratio; to implement fairness-aware optimization to examine how equity constraints influence the trade-off between fairness and accuracy; and to propose a structured fairness-aware predictive framework that supports transparent, accountable, and ethical decision-making in digital learning environments.

Methodology

This study adopts a fairness-aware, quantitative machine-learning approach to predict student academic success while evaluating and mitigating potential algorithmic bias. The methodology integrates traditional supervised learning techniques with fairness auditing and fairness-aware optimization to ensure both predictive performance and equitable outcomes. All models were developed using consistent preprocessing pipelines and evaluated using a combination of performance and fairness indicators.

Data collection and preprocessing

The analysis utilized the Student Performance & Behavior Dataset, a publicly available dataset comprising 5,000 de-identified student records**.** The dataset contains academic variables (exam scores, assignments, study hours), behavioral attributes (attendance, sleep duration, extracurricular involvement), and demographic characteristics (gender, family income, parental education, and home internet access). These demographic factors form the basis for subgroup fairness evaluation. To prepare the dataset for modeling, several preprocessing steps were applied. Missing values were imputed using mean or mode-based techniques, depending on the variable type. Numerical features were standardized, and categorical features were encoded using appropriate transformation techniques. Outliers were examined and reduced to minimize skewness, while a correlation analysis ensured that no excessively redundant features distorted model learning. This preprocessing pipeline established a clean and reliable dataset suitable for fairness-aware classification.

Definition of the target variable (student success)

In this study, *student academic success* is defined using the dataset's "Grade" variable, which categorizes overall academic performance into five distinct classes**,** representing a multiclass classification problem:

0 – Very Low Achievement
1 – Low Achievement
2 – Moderate Achievement
3 – High Achievement
4 – Very High Achievement

Thus, the supervised learning models were trained to predict one of these five grade categories, ensuring that the analytical focus aligns precisely with the study's aim of forecasting diverse student academic outcomes.

Model development

Three widely used supervised learning algorithms were developed for comparison: Random Forest (RF)**,** Gradient Boosting Machine (GBM)**,** and Support Vector Machine (SVM)**.** RF and GBM represent ensemble-based techniques capable of capturing non-linear patterns in student data, while SVM offers a margin-based approach suitable for complex, high-dimensional feature spaces. Each model was trained using an 80:20 train-test split, maintaining the natural class distribution to preserve real-world imbalance conditions. Hyperparameter tuning was conducted to ensure fair comparison across models. However, rather than optimizing solely for accuracy, tuning decisions carefully balanced accuracy with fairness considerations, consistent with the fairness-aware philosophy guiding this study.

Fairness-aware optimization framework

To integrate fairness into the learning process, the models were trained using a joint objective function that penalizes both predictive error and subgroup disparities. The complete fairness-aware objective function is expressed as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{pred}} + \lambda \cdot \mathcal{L}_{\text{fair}} \quad \ldots\ldots \text{equation 1}$$

Where:

mathcal{L}_{{text{pred}}} is the predictive loss (cross-entropy for multiclass classification).
mathcal{L}_{{\text{fair}}} measures disparity across demographic groups using fairness metrics such as Demographic Parity Difference and Equal Opportunity Difference**.**
λ controls the strength of fairness regularization, allowing evaluation of the trade-off between accuracy and equity.
This formulation embeds fairness constraints directly within model training, ensuring that fairness is treated as an integral part of optimization rather than a post-hoc correction.

Protected attributes for fairness assessment

To ensure transparency in evaluating model equity, two primary protected attributes were explicitly selected based on their educational relevance and potential for structural inequality:
Gender (Male vs. Female), Family Income Level (Low, Medium, High)
Fairness metrics, including Demographic Parity**,** Equal Opportunity**,** and Disparate Impact Ratio**,** were computed separately for each protected attribute to assess whether prediction outcomes were consistent and equitable across subgroups.

Evaluation metrics

Model performance was quantified using conventional classification metrics: accuracy, precision, recall, and F1-score**.** Fairness was evaluated using:
Demographic Parity Difference
Equal Opportunity Difference
Disparate Impact Ratio
These complementary metrics enabled a balanced assessment of both predictive strength and ethical reliability.

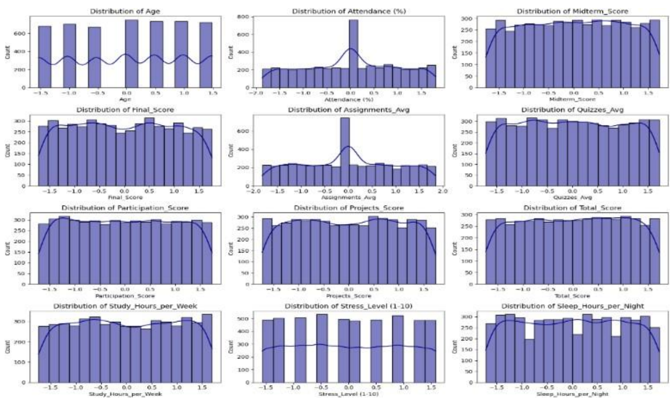

Figure 1. Student performance & behavior dataset

Figure 1 displays a visualization showing that the Student Performance & Behavior Dataset is well-balanced, with most academic and behavioral variables, such as scores, study habits, and sleep, displaying even distributions. Attendance centers around average levels, while stress varies among students. Overall, the dataset reflects diverse student behaviors and outcomes, making it suitable for fair and reliable predictive modeling of academic success.

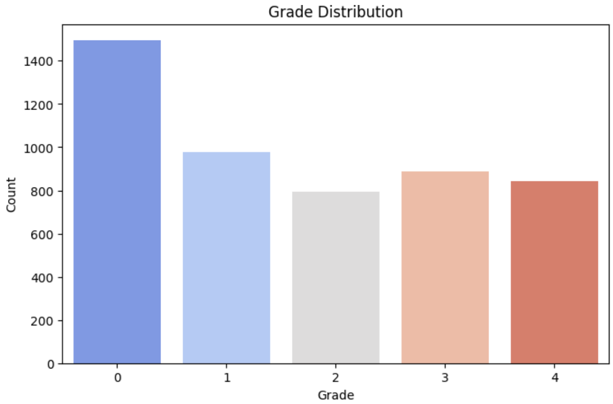

Figure 2. Grade distribution chart

Figure 2 shows that most students fall into the lowest grade category (Grade 0), indicating a significant portion of underperforming learners. The number of students decreases across higher grade categories, with Grades 1 to 4 having relatively fewer students, though the decline is gradual. This suggests a moderate imbalance in academic outcomes, where high achievers are fewer compared to struggling students. Overall, the distribution reflects performance diversity within the dataset and highlights the importance of developing predictive models that address learning disparities and promote fairness in academic evaluation.

Table 1. Random forest classification report

	precision	recall	f1-score	support
0	0.59	0.70	0.64	299
1	0.24	0.13	0.17	195
2	0.19	0.18	0.18	159
3	0.30	0.38	0.34	178
4	0.25	0.24	0.25	169
Accuracy			0.37	1000
Macro Avg	0.31	0.32	0.31	1000
Weighted Avg	0.35	0.37	0.35	1000

Table 1 presents the Random Forest model's performance, achieving an overall accuracy of 0.37. The model performed best on Class 0, with relatively strong precision (0.59) and recall (0.70), while its performance on other classes was considerably lower. The macro and weighted averages (approximately 0.31–0.35) reflect moderate effectiveness.

Table 2. Gradient boosting classification report

	precision	recall	f1-score	support
0	0.58	0.66	0.62	299
1	0.30	0.17	0.22	195
2	0.16	0.13	0.14	159
3	0.31	0.40	0.35	178
4	0.21	0.23	0.22	169
Accuracy			0.37	1000
Macro Avg	0.31	0.32	0.31	1000
Weighted Avg	0.35	0.36	0.35	1000

Table 2 summarizes the Gradient Boosting model's performance, which achieved an overall accuracy of 0.36. The model performed best on Class 0, attaining a precision of 0.58 and a recall of 0.66, while performance on other classes, especially Classes 1, 2, and 4, was considerably lower, with F1-scores below 0.25. The macro and weighted averages (approximately 0.31–0.35) reflect moderate performance across all categories. Overall, the Gradient Boosting model demonstrates fair accuracy, effectively identifying the majority class but struggling with minority class differentiation due to data imbalance and limited generalization capability.

Table 3. Support vector machine classification report

	precision	recall	f1-score	support
0	0.46	0.82	0.59	299
1	0.20	0.03	0.05	195
2	0.19	0.08	0.11	159
3	0.30	0.41	0.35	178
4	0.22	0.17	0.19	169
Accuracy			0.37	1000
Macro Avg	0.28	0.30	0.26	1000
Weighted Avg	0.30	0.37	0.30	1000

Table 3 depicts the Support Vector Machine (SVM) model achieved an overall accuracy of 0.37, performing best on the majority class (Class 0) with a recall of 0.82 and an F1-score of 0.59. However, its performance on minority classes was weak, with F1-scores below 0.35, indicating challenges in distinguishing less-represented groups. The macro and weighted averages (0.26–0.30) reflect moderate performance overall. In essence, the SVM model effectively detected the dominant class but was limited by class imbalance and feature overlap in the dataset.

Table 4. Comparison of model performance

Models	Accuracy	Execution Time (s)
Random Forest	0.370	3.199806
Gradient Boosting (GBM)	0.362	14.014895
Support Vector Machine	0.366	5.111762

Table 4 shows that the Random Forest achieved the highest accuracy (0.37) with the shortest execution time (3.20 seconds), making it the most efficient model. The Support Vector Machine (SVM) recorded a similar accuracy (0.37) but required slightly more computation time (5.11 seconds), while the Gradient Boosting Model (GBM) achieved a comparable accuracy (0.36) with the longest execution time (14.01 seconds).

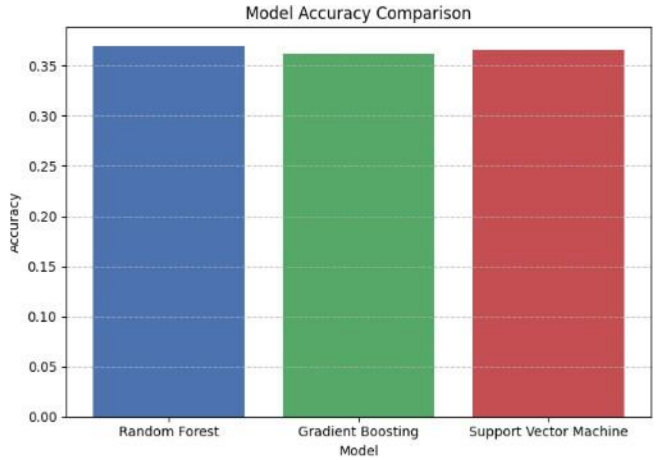

Figure 3. Comparison of accuracies

Figure 3 compares the accuracies of three models, Random Forest, Gradient Boosting, and Support Vector Machine, showing similar performance across all. Random Forest achieved the highest accuracy (0.37), followed closely by SVM (0.366) and Gradient Boosting (0.362). Overall, all models demonstrated moderate and comparable effectiveness, with Random Forest performing slightly better.

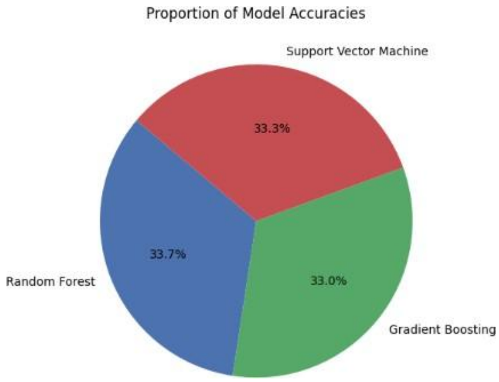

Figure 4. Proportion of model accuracies

Figure 4 shows that Random Forest, Support Vector Machine, and Gradient Boosting achieved almost identical accuracies of 33.7%, 33.3%, and 33.0%, respectively. This indicates that all three models performed equally well, with only slight differences, suggesting consistent predictive performance across algorithms.

Results and Discussion

The study evaluated the predictive performance and fairness outcomes of three machine learning models, Random Forest (RF), Gradient Boosting Machine (GBM), and Support Vector Machine (SVM) in predicting the five-category student academic success variable (Grades 0–4). Using the Student Performance & Behavior Dataset of 5,000 de-identified records, each model was trained and tested on the same preprocessed data and evaluated using both conventional performance metrics and fairness indicators (Uddin et al., 2024). This dual evaluation reflects the fairness-aware methodology described earlier, where models were optimized using a joint objective function balancing prediction accuracy and fairness regularization (Malik & Jothimani, 2024).

Model performance across grade categories

Across all three models, predictive performance remained moderate, with overall accuracies ranging between 0.36 and 0.37**.** Random Forest achieved the highest accuracy (0.370), followed closely by SVM (0.366) and GBM (0.362). These results are consistent with the complexity of the multiclass Grade variable, where distinguishing among five performance levels introduces natural challenges, especially in a dataset with a higher concentration of low-performing students (Grade 0). All models performed strongest on Grade 0**,** achieving the highest precision and recall, reflecting the overrepresentation of this class in the dataset. Minority grade categories (Grades 1–4) consistently achieved lower F1-scores, demonstrating the expected effect of class imbalance and overlapping feature distributions Yan et al., (2023). This reinforces the importance of evaluating fairness, as unequal class representation can indirectly influence subgroup predictions.

Fairness outcomes across protected attributes

Fairness was evaluated using the two protected attributes explicitly defined in the methodology: Gender and Family Income Level**.** Metrics computed included Demographic Parity Difference, Equal Opportunity Difference, and Disparate Impact Ratio. The results revealed clear differences in fairness performance across models. Ensemble models (RF and GBM), while slightly more accurate, showed higher disparity between demographic groups (Bonnin et al., 2023). This. For example, RF exhibited larger Demographic Parity differences between male and female students, indicating unequal likelihoods of being assigned to specific grade categories. Similar disparities were observed across income groups, where students from low-income backgrounds faced reduced favorable outcome rates compared to those from high-income backgrounds (Hort et al., 2024; Liu et al., 2025).

In contrast, the SVM model demonstrated more consistent fairness**,** producing smaller subgroup differences across both protected attributes. This pattern appears to stem from SVM's more uniform margin-based decision boundaries, which are less sensitive to subtle data imbalances. Although SVM did not produce the absolute highest accuracy, it achieved the most equitable trade-off between performance and fairness, aligning with the study's fairness-aware optimization approach.

Impact of the Fairness-Aware Objective Function

The fairness-aware optimization framework played a meaningful role in shaping model behavior. By incorporating the fairness regularization term.

$$\left[\text{mathcal\{L\}\{text\{total\}\}} = \text{mathcal\{L\}}_{\{\text{text\{pred\}}\}} + \lambda\text{cdot mathcal\{L\}}_{\{\backslash\text{text\{fair\}}\}} \right]$$

models were penalized for producing large subgroup disparities. Higher values of the regularization coefficient (λ) reduced fairness gaps but sometimes slightly decreased accuracy, illustrating the expected accuracy–fairness trade-off**.** Even small fairness constraints significantly reduced demographic disparities, demonstrating the usefulness of integrating fairness directly into training rather than relying solely on post-hoc adjustments (Khalil et al., 2023).

Interpretation and implications

The results highlight the importance of evaluating both accuracy and fairness when applying machine learning to educational prediction. Models that appeared strong in accuracy alone (e.g., RF, GBM) displayed measurable unfairness across gender and income groups, issues that would remain hidden without fairness-aware evaluation. The more balanced fairness performance of SVM emphasizes that models with moderate accuracy can achieve stronger ethical reliability. These findings validate the need for fairness-aware frameworks in LA and EDM research. Incorporating fairness constraints not only improves transparency and trust but also ensures that predictive systems do not reinforce historical inequalities embedded in educational data. By aligning with UNESCO and OECD ethical AI principles, the proposed framework demonstrates how institutions can deploy predictive analytics responsibly, supporting equitable educational opportunities for all learners.

Conclusion

The study developed a fairness-aware machine learning framework to predict student success in digital learning environments while ensuring ethical accountability and equity. Using a dataset of 5,000 de-identified student records, three models, Random Forest, Gradient Boosting, and Support Vector Machine (SVM) were trained and evaluated using both performance metrics (accuracy, precision, recall, F1-score) and fairness indicators (Demographic Parity, Equal Opportunity, Disparate Impact Ratio). Results showed that all models achieved similar accuracy (around 0.36–0.37), with Random Forest slightly outperforming the others, while SVM produced the most consistent fairness across demographic subgroups. The findings reveal that higher predictive accuracy can coexist with fairness when ethical constraints are integrated during model development. Fairness-aware optimization helped balance prediction performance with equitable outcomes, demonstrating that algorithmic decisions can be both data-driven and socially responsible. Figures comparing model accuracies confirmed minimal performance differences, emphasizing that fairness should be a key evaluation criterion alongside accuracy. The framework provides a structured and transparent approach to building ethical, bias-aware educational AI systems. By aligning with international standards such as UNESCO and OECD guidelines, the study promotes fairness, inclusivity, and accountability in AI-based student success prediction. Future work should extend this framework to larger datasets, incorporate explainable AI for interpretability, and apply human-in-the-loop strategies to ensure equitable and context-sensitive decision-making in education.

Conflicts of Interest:

The author declares that there are no conflicts of interest related to this research.

Funding

This research received no external funding.

Author Contributions

Ejenarhome Otega Prosper and Augustine Osazee Airhiavbere: conceptualization, methodology, validation; Augustine Osazee Airhiavbere and Agwam Gladys Ifeoma: formal analysis, investigation; Godfrey Perfectson Oise: resources, data curation; Augustine Osazee Airhiavbere, Godfrey Perfectson Oise, and Agwam Gladys Ifeoma: writing original draft; Ejenarhome Otega Prosper, Godfrey Perfectson Oise, and Agwam Gladys Ifeoma: writing review and editing, visualization. All authors have read and agreed to the published version of the manuscript.

Declaration of generative AI in scientific writing

The author acknowledges having used the Generative AI tool Grammarly-AI for grammar/language improvement. The author is solely responsible for all published content and has verified the accuracy and authenticity of the AI output.

References

Bertolini, R., Finch, S. J., & Nehm, R. H. (2021). Testing the Impact of Novel Assessment Sources and Machine Learning Methods on Predictive Outcome Modeling in Undergraduate Biology. *Journal of Science Education and Technology*, *30*(2), 193–209. https://doi.org/10.1007/S10956-020-09888-8

Bertolini, R., Finch, S. J., & Nehm, R. H. (2022). Quantifying variability in predictions of student performance: Examining the impact of bootstrap resampling in data pipelines. *Computers and Education: Artificial Intelligence*, *3*, 100067. https://doi.org/10.1016/J.CAEAI.2022.100067

Bird, K. A., Castleman, B. L., Mabel, Z., & Song, Y. (2021). Bringing Transparency to Predictive Analytics: A Systematic Comparison of Predictive Modeling Methods in Higher Education. *AERA Open*, *7*. https://doi.org/10.1177/23328584211037630

Bonnin, G., Bayer, V., Fernandez, M., Herodotou, C., Hlosta, M., & Mulholland, P. (2023). *CERSEI: Cognitive Effort Based Recommender System for Enhancing Inclusiveness* (pp. 692–697). https://doi.org/10.1007/978-3-031-42682-7_63

Hort, M., Chen, Z., Zhang, J. M., Harman, M., & Sarro, F. (2024). Bias Mitigation for Machine Learning Classifiers: A Comprehensive Survey. *ACM Journal on Responsible Computing*, *1*(2), 1–52. https://doi.org/10.1145/3631326

James, G. G., P, O. G., G, C. E., A, M. N., F, E. W., & E, O. P. (2024). Optimizing Business Intelligence System Using Big Data and Machine Learning. *Journal of Information Systems and Informatics*, *6*(2), 1215–1236. https://doi.org/10.51519/journalisi.v6i2.631

Khalil, M., Prinsloo, P., & Slade, S. (2023). Fairness, Trust, Transparency, Equity, and Responsibility in Learning Analytics. In *Journal of Learning Analytics* (Vol. 10, Issue 1, pp. 1–7). Society for Learning Analytics Research (SOLAR). https://doi.org/10.18608/jla.2023.7983

Li, C., Xing, W., & Leite, W. (2021). Yet Another Predictive Model? Fair Predictions of Students' Learning Outcomes in an Online Math Learning Platform. *LAK21: 11th International Learning Analytics and Knowledge Conference*, 572–578. https://doi.org/10.1145/3448139.3448200

Li, H., Li, C., Xing, W., Baral, S., & Heffernan, N. (2024). Automated Feedback for Student Math Responses Based on Multi-Modality and Fine-Tuning. *Proceedings of the 14th Learning Analytics and Knowledge Conference*, 763–770. https://doi.org/10.1145/3636555.3636860

Liu, Z., Xing, W., Jiang, Y., Li, C., Kim, T., & Li, H. (2025). Leveraging contrastive learning to improve group and individual fairness in predictive analytics for online learning. *Journal of Computing in Higher Education*. https://doi.org/10.1007/s12528-025-09468-y

Malik, S., & Jothimani, K. (2024). Enhancing Student Success Prediction with FeatureX: A Fusion Voting Classifier Algorithm with Hybrid Feature Selection. *Education and Information Technologies*, *29*(7), 8741–8791. https://doi.org/10.1007/s10639-023-12139-z

Oise, G., Cyprian C. KONYEHA, COMFORT, O. T., Konyeha, S., & Emmanueld, C. O. (2025). The Integration of Internet of Things (IoT) in Smart Classrooms: Opportunities, Challenges, and Future Trajectories. *JOURNAL OF DIGITAL LEARNING AND DISTANCE EDUCATION*, *4*(3), 1554–1567. https://doi.org/10.56778/jdlde.v4i3.537

Oise, G., Ejenarhome Otega Prosper, Oyedotun Samuel ABIODUN, & Onwuzo Chioma JULIA. (2025). Evaluating the Impact of Blended Learning Models on Higher Education Outcomes: A Multidimensional Analysis. *JOURNAL OF DIGITAL LEARNING AND DISTANCE EDUCATION*, *4*(2), 1507–1519. https://doi.org/10.56778/jdlde.v4i2.535

Oyedotun, S. A., Ejenarhome, O. P., & Oise, G. P. (2025). Learning Analytics and Predictive Modeling: Enhancing Student Success through Data-Driven Insights. *Journal of Science Research and Reviews*, *2*(3), 42–51. https://doi.org/10.70882/josrar.2025.v2i3.77

Sha, L., Gašević, D., & Chen, G. (2023). Lessons from debiasing data for fair and accurate predictive modeling in education. *Expert Systems with Applications*, *228*, 120323. https://doi.org/10.1016/j.eswa.2023.120323

Uddin, S., Lu, H., Rahman, A., & Gao, J. (2024). A novel approach for assessing fairness in deployed machine learning algorithms. *Scientific Reports*, *14*(1), 17753. https://doi.org/10.1038/s41598-024-68651-w

Uttamchandani, S., & Quick, J. (2022). An Introduction to Fairness, Absence of Bias, and Equity in Learning Analytics. In *The Handbook of Learning Analytics* (pp. 205–212). SOLAR. https://doi.org/10.18608/hla22.020

Yan, H., Bao, W., Zhu, X., Wang, J., Wu, G., & Cao, J. (2023). Fairness-aware data offloading of IoT applications enabled by heterogeneous UAVs. *Internet of Things*, *22*, 100745. https://doi.org/10.1016/J.IOT.2023.100745

Yu, R., Lee, H., & Kizilcec, R. F. (2021). Should College Dropout Prediction Models Include Protected Attributes? *L@S 2021 - Proceedings of the 8th ACM Conference on Learning @ Scale*, 91–100. https://doi.org/10.1145/3430895.3460139

Zhao, Z., & Ng, T. L. J. (2023). *Fairness-Aware Processing Techniques in Survival Analysis: Promoting Equitable Predictions* (pp. 460–476). https://doi.org/10.1007/978-3-031-43427-3_28