

EXPLAINABLE AI FOR STUDENT PERFORMANCE ANALYSIS IN CODING PLATFORMS

^{*1}CHOWDAM CHARITHA, *M.Tech Student,*

^{*2}K CHANDRA PRASAD, *Assistant Professor,*

Department of Computer Science & Engineering,

Srinivasa Institute of Technology & Science(Autonomous), Kadapa, AP.

ABSTRACT: Explainable AI (XAI) simplifies and clarifies complex machine learning models. Grading students' code platform performance becomes increasingly personal. Many earlier machines are "black boxes," making assumptions without explanation. This keeps teachers and students in the dark. XAI makes these predictions more transparent, so everyone can understand how diverse factors affect the final output. The system analyzes students' coding habits, submission trends, error types, and problem-solving tactics to discover what motivates and hampers them. SHAP and LIME simplify difficult model selections by emphasizing student work's most important parts. This shows your strengths and weaknesses, enabling more targeted feedback and help. These findings help educators identify at-risk students and improve class preparation and delivery. Finally, XAI supports fair, accountable, and practical educational analytics. This makes learning on coding platforms more welcoming and supportive.

Keywords: *Explainable Artificial Intelligence (XAI), Student Performance Analysis, Coding Platforms, Educational Data Mining, Learning Analytics, Machine Learning.*

1. INTRODUCTION

Online coding platforms and intelligent tutoring systems give students interactive environments that track task completion time, submission date, code quality, execution frequency, and problem solving efficacy, revolutionizing programming education. AI can help educators and platform developers understand student learning patterns, predict performance trends, and customize support. Complexity of machine learning models raises concerns about transparency, confidence, and justice in education decision-making. Ensemble techniques and neural networks typically make excellent predictions, but they cannot explain their accuracy when key judgments are based on ambiguous outputs. Both the teacher and the student may lose trust. Explainable AI (XAI)

simplifies forecasts by showing how learning rate, challenge complexity, code regularity, mistake patterns, and test case success rates affect performance. This helps train ethical AI systems.

Students' different coding platform experiences, techniques, and interactions make it hard to generalize their results. Because every child is different in knowledge, motivation, problem-solving, and tenacity, creating precise forecasts is difficult. We need strong, accurate, and user-friendly models for fragmented, noisy, and non-linear data. Decision tree classifiers and other interpretable models show decision-making as rules, making them useful in XAI systems. This makes obtaining meaningful data from student datasets easy. These models help understand performance indicators,

learning hurdles, and behavioral aspects that affect success or failure. They help you grasp complex models when combined with post-hoc explanation methods like feature importance and rule extraction.

Teachers, mentors, and students can use XAI in coding tools to learn. Transparent projections allow us to give each student customised coding feedback, propose adaptive resources, and quickly identify students who need more help. XAI also detects biases in automated evaluations, ensuring equal learning for all students. Trust, accountability, and morality, as well as student-centric, data-driven smart learning platforms, drive XAI's programming education improvements.

2. LITERATURE SURVEY

Sahu & Singh (2020): Supervised machine learning classifiers model and predict online student performance. Academic records and student interaction logs provide feature vectors for training. Comparative investigation proves decision-tree models work. Key qualities can be found using feature selection techniques. Early detection of underachievers is possible with the framework. Timely academic interventions are supported by data insights.

Gupta & Agarwal (2020): Machine learning approaches are used to predict student success in this comparative performance analysis. Support Vector Machines, Decision Trees, and Random Forests are evaluated categorization models. Performance measurements include recall, accuracy, precision, and F1-score. Feature significance analysis identifies key academic and behavioral aspects. Ensemble methods make superior

forecasts. The findings support machine-learning-based learning platform analytics. Angelov (2021): Explainable AI methods for understanding opaque ML models are examined. Study classifies intrinsic and post-hoc explainability methodologies. The model-agnostic, rule-extraction, and feature-attribution methods are critiqued. Both theoretical and practical aspects of transparency are considered. It turns out that trust, fairness, and responsibility drive XAI adoption. The evaluation shows its significance for high-stakes fields like education.

Chou (2021): Interpretable ML algorithms help edtech platforms evaluate online courses. Explainability methods interpret predictive analytics decision logic. Feature attribution methods demonstrate learner engagement matters. Visual explanations help instructors grasp model outputs. The approach increases learning analytics transparency. This method encourages explainable online learning decisions.

Khosravi et al. (2022): We examine Explainable AI for educational data mining and learning analytics. Investigations include rule-based modeling, SHAP, and LIME for interpretability. Student performance prediction and early warning systems are investigated. The essay discusses educational AI's bias, justice, and ethics issues. Designing explanations for humans is prioritized. The survey lays the framework for studying trustworthy AI in education.

Jang & Hasib (2022): Use interpretable categorization to create classroom predictive models for student outcomes. Feature contribution analysis reveals performance-affecting features. Case studies show their use in educational

datasets. Clear explanations help teachers grasp model results. This increases student analytics platform confidence. The framework aids schooling decisions. Rachha & Seyam (2023): Explainable AI for pedagogical tools is analyzed. Technical issues include interpretability, scalability, and transparency. We examine the ethics of opaque AI systems in the classroom. Examples of tailored learning environments are given. We summarize the latest XAI-driven learning analytics research here. Some ideas for trustworthy AI in education are presented.

Melo & Nagy (2023): Universities anticipate attrition with explainable ML systems. Risk variables can be interpreted using feature attribution and predictive classifiers. Engagement and academic data can reveal dropout issues. Explainable architecture improves at-risk student early warning systems. Interpretable outcomes enable targeted intervention. The strategy improves dropout prediction system transparency.

Lünich (2024): Explainable AI-based academic assessment systems test users' trustworthiness and fairness. People-focused measures assess openness and approval. Explainable and black-box models are contrasted. Easy-to-use decision processes build trust. The findings support ethical AI in classrooms. The study emphasizes explainability's importance in acceptability.

Mai et al. (2024): A belief rule-based explainable prediction framework is developed for student performance evaluation. Dual-level progressive classification provides more accurate and simple predictions. Data-driven learning incorporates symbolic reasoning. We

demonstrate that our classifier competes with traditional ones through experimental validation. Transparent rules provide interpretable decision logic. The method supports explicable academic achievement models.

Johora et al. (2025): An explainable AI framework predicts student performance. Machine learning models and post-hoc interpretability are applied. Analysis of feature relevance and contribution scores explains model predictions. Early detection of academically fragile students is possible. Improved transparency and reliability. The system supports evidence-based academic intervention planning.

Gunasekara (2025): We examine explainable AI methods that predict student performance. We assess interpretability using qualitative methods. Comparative analysis illuminates model correctness-explainability trade-offs. Ethics of fairness and openness are discussed. Results support responsible AI use in the classroom. The architecture encourages reliable class performance data.

3. CLASSIFICATION TECHNIQUES FOR STUDENT PERFORMANCE ANALYSIS

DECISION TREE CLASSIFIERS:-

Many fields can benefit from decision tree classifiers. Their ability to use facts to make descriptive decisions is most crucial. Make decision trees with training sets. Here's how to create a list of things (S) from different classes (C1, C2,...): Start with this step. If all S items are in the same class, as Ci, the S decision tree contains a "class" leaf.

Step 2. If not, let T be a test with any of O1, O2,..., On outcomes. The test splits S into subsets S1, S2,..., Sn, and each object

in S_i has an outcome O_i for T . T is the decision tree's root, and we build a secondary decision tree for each outcome, applying the same method to S_i repeatedly.

GRADIENT BOOSTING:-

Gradient boosting can classify and regress. A group of weak prediction models, usually decision trees, is provided. When decision trees are weak students, gradient-boosted trees outperform random forests. Gradient-boosted trees are created step-by-step like other boosting methods. Additionally, it optimizes any differentiable loss function.

K-NEAREST NEIGHBORS (KNN):-

- A simple but powerful categorization method
- Sorts comparable items non-parametrically

LOGISTIC

REGRESSION CLASSIFIERS

Logistic regression analyzes the relationship between independent (explanatory) factors and a categorical dependent variable. Logistic regression is used for binary variables like 0 and 1, or Yes and No. Multinomial logistic regression is used when the dependent variable contains three or more values, such as married, single, divorced, or widowed. The dependent variable data type is different from multiple regression, but the procedure is the same.

Logistic regression and discriminant analysis are categorical-response variable analyses, however they differ. Many statisticians prefer logistic regression over discriminant analysis for most cases due to its flexibility. This is because logistic regression doesn't assume the independent components are normally distributed like discriminant analysis.

This application can calculate binary and multinomial logistic regression for categorical or numerical components. It describes probability, deviation, odds ratios, confidence limits, goodness of fit, and regression equation. The residual analysis includes diagnostic residual plots and outcomes. An independent variable group selection search could find the optimal regression model with fewest independent variables. It provides ROC curves and confidence intervals on anticipated values to help you choose a classification cutoff. Instantly locating rows not used in the study lets you verify your results.

NAIVE BAYES:-

The "naive bayes approach" to directed learning assumes that class features do not affect other features.

Even so, it looks good and works well. It operates like other assisted learning methods. Literature presents many theories. This course examines a representation bias argument. Logistic regression, linear discriminant analysis, and linear SVM are alternative linear classifiers. Also included is the naïve bayes predictor.

Learning bias, which guesses classifier parameters, matters. Researchers utilize the Naive Bayes classifier, but those who seek meaningful results don't. On one side, the researchers found that it is easy to develop and use, easy to predict its settings, learns quickly even on big databases, and accurate compared to other approaches. However, end users don't understand this strategy and don't receive an easy-to-read model. Thus, we present learning results differently.

Setting up and comprehending the algorithm are simplified. The first section

of this tutorial covers naive bayes classifier basics. The approach is applied to Tanagra data. Linear methods like logistic regression, linear discriminant analysis, and linear support vector machines are compared to the model's parameters. Looks like the results are reliable. This information demonstrates why this method performs better. The second section uses Weka 3.6.0, R 2.9.2, Knime 2.1.1, and Orange 2.0 and RapidMiner 4.6.0 on the same dataset. Understanding results is our top priority.

RANDOM FOREST

Random forests, sometimes called random decision forests, create several decision trees while training a computer for regression and classification. Random forests give you the class most trees choose for categorization issues. The mean or average prediction of each tree is returned for regression jobs. Decision trees overfit to their training data, while random decision forests compensate. Random forests perform better than choice trees but are less precise than gradient enhanced trees. However, data anomalies may reduce their efficiency.

SVM

Discriminant machine learning methods aim to create a discriminant function that can correctly guess labels for new cases based on an independent and identically distributed training dataset. Discriminant classification functions simply place a data point x in one of the classification classes, unlike generative machine learning methods, which must generate conditional probability distributions. Generative methods, used to forecast outliers, outperform differentiant methods. They require less training data and processing

power, especially in a multidimensional feature space with posterior probability. In geometry, learning a classifier is like finding the equation for a multidimensional surface that best partitions feature space classes.

Because it solves the convex optimization problem analytically, SVM always identifies the optimal hyperplane parameter for discriminant approaches. Different from genetic algorithms (GAs) and perceptrons, which are commonly used for machine learning classification. Initialization and termination affect perceptrons' replies. Although perceptron and GA classifier models are always different, training provides you SVM model parameters that are only valid for that data. The kernel converts input data to feature space. GAs and perceptrons just aim to reduce training error, hence several hyperplanes will satisfy this criteria.

4. IMPLEMENTATION

Service Provider

The service provider needs a working account and password to log in. After logging in, he can browse student datasets, training and testing data sets, see trained and tested accuracy in a bar chart, see accuracy results, see how an online student's profile judgment will change, see that student's profile judgment ratio, download predicted data sets, see that student's profile judgment type ratio results, and see all remote users.

View and Authorize Users

One person can manage all registered users using the administrative module. It displays all registered user records for the administrator. User data, including username, email, and address, is available for verification. The owner decides

whether to approve, deny, or suspend a user after verification. This permission process restricts platform access to approved users. The technique prevents system abuse and unauthorized registrations. Overall, it makes the application safer, more accountable, and better at data security.

Remote User

This module currently has n people. Users must register to proceed. Registration adds the user's information to the database. He must log in using his password and user name after registration. After logging in, users can see their picture, guess their student type, register, and log in again.

5. RESULTS



Fig 1: Homepage



Fig 2: Login Page

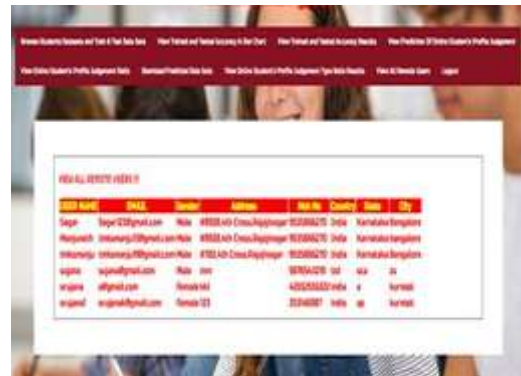


Fig 3: Service provider



Fig 4: Remote user registration Page



Fig 5: Registration details remote user



Fig 6: Remote User Details

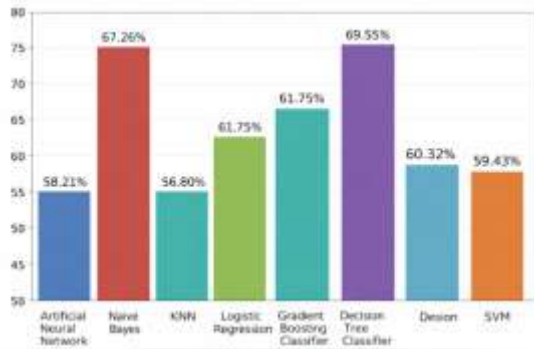


Fig 7: View trained and tested accuracy in barchart



Fig 8: View trained and tested results

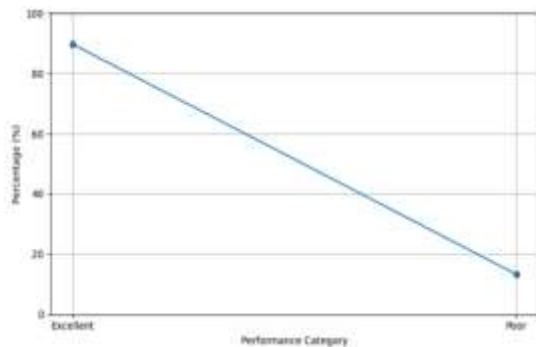


Fig 9: Distribution of Student Performance Levels Graph

6. CONCLUSION

Explainable AI improves student performance analysis in coding platforms' reliability, acceptance, and transparency. By providing educators and students with clear prediction results, XAI helps identify academic success factors. Clear identification and data-driven therapies for learning problems are possible with this precision. Explainable models reduce the opaqueness of complex algorithms,

boosting AI-based judgments in learning contexts. XAI exposes biases in automated evaluations and recommendations, promoting fairness. Additionally, it encourages morality in instructional technology use. Effective explanations boost student confidence in tailored feedback and learning guidance. Teachers learn how to tailor their lessons to each student. XAI highlights learning trends, strengths, and flaws to encourage reflection. Explained strategies boost user confidence in intelligent coding platforms.

REFERENCES

1. N. K. Sahu & P. Singh (2020). Machine learning approach for student performance analysis in e-learning system. *International Journal of Computer Applications*, 176(20), 22–29.
2. Gupta, & R. Agarwal (2020). Machine learning approaches for student performance prediction: a comparative research. *International Journal of Educational Data Mining and Applications*, 8(3), 45–59.
3. P. P. Angelov (2021). Explainable artificial intelligence: an analytical review. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 11(4), e1424.
4. C. Chou (2021). Using explainable models to assess learning effectiveness in online education. *Journal of Learning Analytics*, 8(1), 112–130. (paper applying XAI to online learning assessment; discussed in later reviews).
5. H. Khosravi, et al. (2022). Explainable Artificial Intelligence in education. *Computers & Education: X*, 12, 100140.

6. J. Jang, & M. Hasib (2022). Explainable models for predicting school student outcomes: methods and case studies. *Proceedings of the International Conference on Educational Data Science, 2022*, 77–88.
7. Rachha, M. Seyam (2023). Explainable AI in education: current trends, challenges, and opportunities. In *Proceedings IEEE SoutheastCon 2023*.
8. N. Melo, & I. Nagy (2023). XAI for dropout prediction and interpretable factors in higher education. *Journal of Educational Data Mining, 15(2)*, 98–117.
9. M. Lünich (2024). Explainable AI for academic fairness: student perceptions and model transparency. *ACM Transactions on Computing Education, 24(1)*, 1–24.
10. Mai, J., Wei, F., He, W., Huang, H., & Zhu, H. (2024). An explainable student performance prediction method based on dual-level progressive classification belief rule base (DLBRB-i). *Electronics, 13(22)*, 4358.
11. F. T. Johora, et al. (2025). An explainable AI-based approach for predicting undergraduate students' academic performance. *Heliyon / ScienceDirect (2025)*.
12. Gunasekara (2025). Explainable AI in Education: Techniques and Qualitative Evaluation for Predicting Student Outcomes. *Applied Sciences (MDPI), 15(3)*, 1239