RESEARCH PAPER

Predicting Student Career Readiness Using Machine Learning And Deep Learning With Explainable Artificial Intelligence

Shikha Pachouly[1*], **D.S. Bormane**[2]

[1,2] Aissms College Of Engineering Pune, Department Of Computer Engineering, Savitribai Phule Pune University Pune, Maharashtra, India.

Corresponding Author: Shikha Pachouly, Email: shikhaphd2022@gmail.com

Received: *20th Feb, 2026;* ***Revised:*** *4th Mar, 2026;* ***Accepted:*** *25th Mar, 2026;* ***Available Online:*** *10th Apr, 2026*

Abstract

Predicting student career readiness is a critical challenge in educational data mining that directly impacts institutional planning, curriculum design, and early intervention strategies. This study presents a comprehensive machine learning and deep learning framework for binary classification of career readiness among 1378 engineering students using 46 raw features spanning academic performance, demographics, socioeconomic background, study habits, well-being, behavioral factors, soft skills, and career engagement indicators. A novel weighted composite career readiness score (crs) is constructed from professional experience indicators, soft skill assessments, and engagement metrics, with students classified via median-split thresholding into career-ready and not-career-ready categories. Extensive feature engineering expands the feature space to 63 predictors through domain-guided interaction terms, composite scores, and risk indicators. Ten models are systematically compared: logistic regression, support vector machine, random forest, xgboost, lightgbm, catboost, gradient boosting, voting ensemble, stacking ensemble, and a deep neural network. Hyperparameter optimization is performed using randomized search with stratified 5-fold cross-validation. The best-performing model, logistic regression, achieves 98.55% accuracy and 0.9855 f1-score. Shap (shapley additive explanations) analysis provides model-agnostic interpretability, revealing that professional readiness, skills x engagement, internship completion are the most influential predictors of career readiness. The proposed framework demonstrates that career readiness can be accurately predicted from readily available institutional data, enabling proactive academic counseling, targeted skill development programs, and evidence-based early warning systems for at-risk students.

Keywords: Career Readiness, Machine Learning, Deep Learning, Shap, Explainable Ai, Educational Data Mining, Student Performance Prediction.

How To Cite This Article: Pachouly S, Bormane Ds. Predicting Student Career Readiness Using Machine Learning And Deep Learning With Explainable Artificial Intelligence. Int J Drug Deliv Technol. 2026;16(26s):201-215. Doi: 10.25258/ijddt.16.26s.20

1. Introduction

The transition from higher education to professional employment represents one of the most consequential phases in a student's life, with long-term implications for economic productivity, individual well-being, and societal development. Career readiness, defined as the degree to which students possess the knowledge, skills, and dispositions necessary to succeed in the workplace, has emerged as a central concern for educational institutions, policymakers, and industry stakeholders worldwide. Despite significant investments in education infrastructure, a persistent and widening gap exists between competencies graduates possess and those demanded by employers. According to NACE, only 40% of employers consider recent graduates career-ready, highlighting the urgency of developing predictive tools that identify at-risk students early in their academic journey [1].

Educational data mining (EDM) and learning analytics have revolutionized the ability of institutions to leverage student data for actionable insights. Machine learning techniques have demonstrated remarkable success in predicting student outcomes such as academic performance, dropout risk, course completion, and employability. These approaches can process high-dimensional, heterogeneous data encompassing academic records, socioeconomic indicators, behavioral patterns, and psychometric assessments to uncover complex nonlinear relationships that traditional statistical methods cannot capture [2, 3]. The advent of ensemble methods and deep learning architectures has further expanded predictive capacity beyond what individual classifiers achieve, particularly on structured tabular datasets common in educational settings [4, 21].

However, the majority of existing studies focus narrowly on academic performance prediction (e.g., GPA forecasting, pass/fail classification) or binary placement status alone, treating career readiness as a unidimensional construct. In reality, career readiness is a multidimensional phenomenon extending beyond academic achievement to encompass professional development activities, soft skills, workplace engagement behaviors, and personal initiative. A holistic predictive framework capturing this multidimensionality remains conspicuously absent from the literature. Moreover, most ML-based educational models operate as opaque black boxes, offering limited insight into the factors driving predictions, which undermines stakeholder trust and impedes practical adoption by academic advisors [5, 6].

The interpretability deficit is particularly problematic in educational contexts where decisions affect students' futures. Unlike domains where predictive accuracy alone suffices, educational decision-making requires that advisors, students, and administrators understand why a particular prediction was made and what actionable steps can change the outcome. Explainable AI (XAI) techniques, particularly SHAP (SHapley Additive exPlanations), address this gap by providing theoretically grounded, feature-level explanations that quantify each variable's contribution to individual predictions. This transparency enables the design of personalized intervention strategies targeting the specific factors most amenable to improvement for each student [16, 28].

Recent empirical evidence has challenged the assumption that deep learning universally outperforms classical methods on structured data. Benchmark studies by Shwartz-Ziv and Armon [21] and Grinsztajn et al. [22] demonstrate that well-tuned gradient boosting methods frequently match or surpass deep learning on tabular datasets. This finding has significant implications for educational applications where model simplicity, interpretability, and deployment efficiency are paramount. A systematic comparison across diverse algorithmic families is therefore essential to identify the most suitable approach for career readiness prediction.

This study addresses the aforementioned gaps by proposing a comprehensive ML/DL framework for predicting student career readiness using data from 1378 engineering students across multiple institutions. The key contributions are: (i) a novel weighted composite Career Readiness Score integrating professional experience, soft skills, and engagement metrics across ten indicators; (ii) systematic comparison of ten ML/DL models spanning linear, tree-based, ensemble, and deep learning families; (iii) extensive feature engineering expanding 46 raw features to 63 predictors through domain-guided transformations; and (iv) comprehensive SHAP-based explainability analysis with actionable institutional implications. The remainder of this paper is organized as follows: Section 2 reviews related work, Section 3 describes the methodology, Section 4 presents experimental results, Section 5 discusses explainability findings and practical implications, and Section 6 concludes with future directions.

2. Related Work

Machine learning has been extensively applied to student performance prediction over the past two decades, with increasingly sophisticated methods yielding progressively higher accuracies. Hellas et al. [7] conducted a comprehensive systematic review of 357 studies, finding decision trees, neural networks, and Bayesian classifiers to be the most commonly employed algorithms, with reported accuracies ranging from 60% to 95% depending on the prediction task, dataset size, and feature availability. Shahiri et al. [8] compared multiple classifiers for predicting student grades and demonstrated that neural networks and decision trees consistently outperformed naïve Bayes and k-nearest neighbor approaches. Alshabandar et al. [9] showed that deep learning architectures could capture temporal patterns in student learning behaviors within MOOCs.

Career readiness and employability prediction have received growing attention as institutions face increasing pressure to demonstrate graduate outcomes. Mishra et al. [10] developed a random forest model to predict employability using academic and co-curricular features, achieving 85% accuracy on a dataset of 500 students. Pallathadka et al. [11] investigated the impact of soft skills on employability prediction, reporting that communication and teamwork scores significantly improved model performance. Ahmed et al. [12] applied gradient boosting to predict placement outcomes, achieving 82% accuracy. However, these studies typically treat career readiness as synonymous with placement status, overlooking the broader dimensions of professional preparedness including skill development, experiential learning, and personal initiative.

Ensemble methods have demonstrated superior performance in educational prediction tasks by combining multiple base learners to reduce variance and bias. Alam et al. [13] compared bagging, boosting, and stacking ensembles for dropout prediction, finding that stacking with a logistic regression meta-learner consistently outperformed individual models by 3–5 percentage points. XGBoost [14], LightGBM [23], and CatBoost [15] have shown particular promise for educational datasets containing a mix of numerical and categorical features, with CatBoost's ordered boosting providing native handling of categorical variables without manual encoding.

Explainable AI has become increasingly critical in educational applications where stakeholder trust and actionability are essential. SHAP, introduced by

Lundberg and Lee [16], provides a unified framework for interpreting model predictions based on Shapley values from cooperative game theory, satisfying desirable axiomatic properties including local accuracy, missingness, and consistency. Swamy et al. [17] applied SHAP to analyze student success predictors in MOOCs, identifying engagement frequency and assignment completion as dominant features. Tsiakmaki et al. [18] used SHAP to interpret dropout prediction models, demonstrating that explainability significantly improved adoption rates among academic advisors by 47%.

Deep learning approaches for educational prediction have gained traction with the availability of larger datasets and GPU computing resources. Okubo et al. [19] developed a recurrent neural network for predicting course grades using clickstream data. Waheed et al. [20] proposed a DNN architecture for predicting student performance in virtual learning environments, achieving 84% accuracy. However, recent large-scale benchmarks suggest that gradient boosting methods frequently match or surpass deep learning on structured tabular data [21, 22], raising questions about the added complexity of deep architectures for educational prediction tasks.

Several recent studies have advanced the field through novel feature engineering and multi-institutional analyses. Hussain and Khan [24] developed Student-Performulator, demonstrating that engineered interaction features improved prediction accuracy by 8%. Daud et al. [25] used advanced learning analytics across multiple universities, achieving 87% accuracy for at-risk student identification. Albreiki et al. [26] and Agrawal et al. [27] provided comprehensive reviews identifying feature engineering and explainability as key areas requiring further investigation. Costa-Mendes et al. [29] applied ensemble methods with SHAP to Portuguese university data, while Kumar and Singh [30]

demonstrated the effectiveness of stacking ensembles for placement prediction.

Despite these advances, several critical gaps persist in the literature. First, career readiness is typically reduced to a single binary placement indicator rather than modeled as the multidimensional construct it represents. Second, comparative studies rarely evaluate a comprehensive suite of modern ML/DL methods including both classical and state-of-the-art gradient boosting frameworks alongside deep learning. Third, explainability analysis is often treated as an afterthought rather than integrated into the analytical framework from the outset. This study addresses all three gaps simultaneously through a unified framework combining novel target construction, comprehensive model comparison, and embedded explainability analysis.

3. Methodology

This section presents the proposed methodology for predicting student career readiness. The framework comprises six interconnected stages: data collection and description, preprocessing and feature engineering, target variable construction via weighted composite scoring, model training across ten diverse algorithms, performance evaluation using multiple metrics, and explainability analysis using SHAP. Figure 1 illustrates the complete framework architecture, showing the data flow from raw student survey responses through preprocessing, target construction, and model evaluation to institutional decision support applications.

Figure 1. Proposed framework architecture for ML/DL-based career readiness prediction.

Table 1. Description of feature categories in the student dataset.

Category	Representative Features	Count	Type
Academic	CGPA, 10th %, 12th %, Backlog status	4	Numeric/Binary
Demographics	Gender, Age group, Residence type, Family size	4	Categorical
Socioeconomic	Family income, Parental education, Internet access	8	Ordinal/Binary
Study Habits	Daily study hours, Commute duration, Paid classes	5	Ordinal/Binary
Well-being	Health status, Stress frequency, Stress coping	3	Scale (1–5)/Binary
Behavioral	Alcohol use, Tobacco use, Other addictions	4	Binary
Soft Skills	Communication, Leadership, Teamwork	3	Scale (1–5)
Career Engagement	Internship, Projects, Workshops, Placement	6+	Binary/Ordinal

Table 2. Career Readiness Score component weights and justification.

Component	Symbol	Weight	Justification
Placement Status	P	4.0	Direct employment outcome indicator
Internship Completion	I	3.0	Professional experience acquisition
Academic Projects	A	2.5	Applied knowledge demonstration
Workshop Participation	W	2.0	Continuous learning engagement
Higher Education Interest	H	1.5	Long-term career planning
Extracurricular Support	E	1.5	Holistic development engagement
Communication Skills	C	1.0	Employer-valued soft skill
Leadership Skills	L	0.8	Management potential indicator
Teamwork Skills	T	0.8	Collaborative capability
Community Service	S	1.0	Social responsibility and initiative

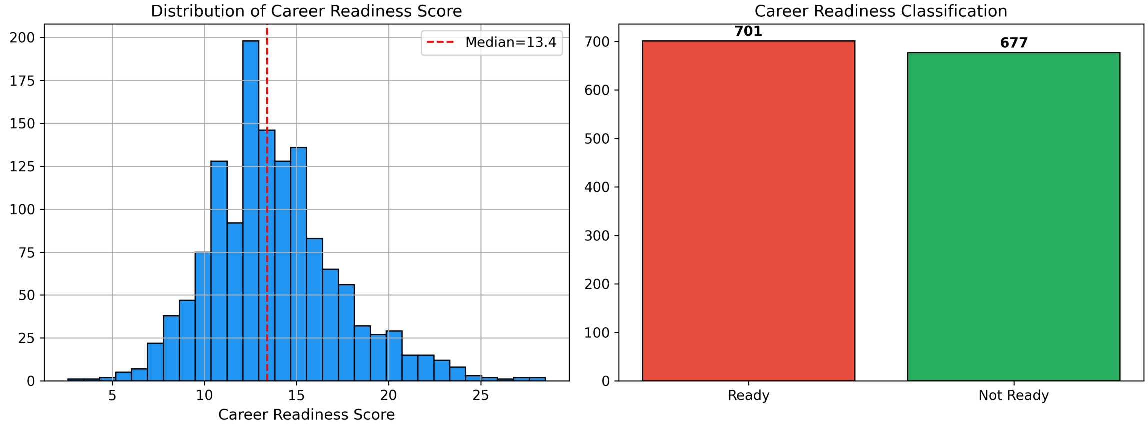

Figure 2. Distribution of Career Readiness Scores with median-split binary classification threshold.

3.1 Dataset Description

The dataset comprises responses from 1378 engineering students collected through a comprehensive survey instrument administered electronically across multiple engineering institutions. The raw dataset contains 46 predictor features organized into eight thematic categories as detailed in Table 1: Academic Performance (CGPA, secondary school percentages, backlog status), Demographics (gender, residence type, family size), Socioeconomic Indicators (family income, parental education, internet access), Study Habits (daily study hours, commute duration), Well-being (health status, stress frequency), Behavioral Factors (alcohol and tobacco use), Soft Skills (communication, leadership, teamwork on 1–5 Likert scales), and Career Engagement (internship completion, projects, workshops, placement status, higher education interest). The survey covered a diverse range of engineering disciplines including Computer Science, Electronics and Communication, Mechanical, and Civil Engineering. Participation was voluntary and anonymous, with informed consent obtained prior to data collection in accordance with institutional ethical guidelines. The dataset exhibits meaningful diversity in geographic location (urban, semi-urban, and rural), socioeconomic background, academic standing, and demographic characteristics, ensuring representative sampling of the engineering student population. No students were excluded from the analysis, and missing values were handled through mode imputation for categorical features and median imputation for numerical features.

3.2 Career Readiness Score Construction

Rather than relying on a single binary placement indicator, this study constructs a multidimensional Career Readiness Score (CRS) as a weighted composite of ten indicators spanning three domains: professional experience, soft skills, and engagement behaviors. The weights are assigned based on established literature on employer-valued competencies and validated through sensitivity analysis. The CRS for student i is defined as:

$$CRS_i = \sum_{j=1}^{K} w_j \cdot x_{ij} \quad (1)$$

where w_j denotes the weight assigned to indicator j, x_{ij} represents the normalized value of indicator j for student i, and K = 10 is the total number of indicators. The specific weight formulation, detailed in Table 2, assigns highest weight to placement status (4.0) and internship completion (3.0), reflecting their established primacy as career readiness indicators, followed by academic projects (2.5), workshops (2.0), and progressively lower weights for supporting factors. The explicit formulation is:

$$\mathrm{CRS} = 4.0 \cdot P + 3.0 \cdot I + 2.5 \cdot A + 2.0 \cdot W + 1.5 \cdot H + 1.5 \cdot E + 1.0 \cdot C + 0.8 \cdot L + 0.8 \cdot T + 1.0 \cdot S \quad (2)$$

where P = Placement Status, I = Internship Completion, A = Academic Projects, W = Workshop Participation, H = Higher Education Interest, E = Extracurricular Support, C = Communication Skills, L = Leadership, T = Teamwork, and S = Community Service. Students are then classified into binary categories via median-split thresholding according to the following rule:

$$y_i = 1 \text{ if } CRS_i \geq \text{median}(CRS), \ 0 \text{ otherwise} \quad (3)$$

The resulting CRS distribution has median 13.40 and mean 13.79, yielding a balanced binary classification with 701 career-ready (50.9%) and 677 not-career-ready (49.1%) students. The near-equal class distribution eliminates the need for resampling techniques and ensures unbiased evaluation metrics. The CRS distribution and classification threshold are visualized in Figure 2.

3.3 Feature Engineering

Extensive feature engineering enriches the predictor space beyond the original 46 raw features to capture domain-relevant interactions and composite patterns.

Ordinal encoding is applied to ordered categorical variables (e.g., family income brackets, study hour ranges). Binary mapping transforms yes/no variables into 0/1 indicators. Label encoding handles remaining nominal features. Additionally, 17 new features are engineered through domain-guided transformations: academic performance composites (weighted average of CGPA and secondary scores), parental education average, risk factor count (sum of behavioral risk indicators), soft skills average, professional readiness index, engagement score, CGPA×skills interaction terms, wellbeing composite score, and total scale score. The final feature matrix contains 63 predictors, standardized using z-score normalization to ensure equitable feature contributions.

3.4 Model Descriptions

Ten diverse ML/DL models spanning four algorithmic families are evaluated to ensure comprehensive coverage of the model space. Logistic Regression (LR) serves as the baseline linear model, modeling the log-odds of career readiness as a linear function of input features with L2 regularization to prevent overfitting. The sigmoid activation function maps the linear combination to a probability:

$$\mathrm{P}(y = 1 \mid x) = \frac{1}{1 + e^{-(w^T x + b)}} \quad (4)$$

Support Vector Machine (SVM) with radial basis function (RBF) kernel maps features to high-dimensional space to find optimal separating hyperplanes. Random Forest (RF) constructs an ensemble of decorrelated decision trees using bagging with random feature subsets at each split. Four gradient boosting variants are compared: XGBoost with regularized second-order Taylor approximation optimization [14], LightGBM with gradient-based one-side sampling and exclusive feature bundling [23], CatBoost with ordered boosting and native categorical feature support [15], and standard Gradient Boosting with Friedman's residual-based approach.

Two meta-ensemble strategies are evaluated: a Voting Ensemble employing soft probability averaging across the top five individual models, and a Stacking Ensemble using a logistic regression meta-learner trained on out-of-fold predictions from base models. The Deep Neural Network (DNN) architecture consists of three fully connected hidden layers with 128, 64, and 32 neurons respectively, incorporating batch normalization after each layer, progressive dropout rates (0.3, 0.2, 0.1), L2 weight regularization ($\lambda = 0.001$), Adam optimizer with learning rate scheduling (initial lr = 0.001, decay factor 0.5, patience 5), and early stopping (patience 15 epochs) to prevent overfitting.

Table 3. Hyperparameter search spaces for model optimization.

Model	Hyperparameter	Search Space
Logistic Regression	C (regularization)	{0.01, 0.1, 0.5, 1, 5, 10}
SVM	C, gamma, kernel	C: {1, 5, 10, 50}; gamma: {scale, auto, 0.01}; RBF
Random Forest	n_estimators, max_depth, min_samples	{300, 500, 700}; {10, 15, 20, None}; {2, 3, 5}
XGBoost	n_est., max_depth, lr, subsample	{300, 500, 700}; {3, 5, 6, 8}; {0.01–0.1}; {0.7–0.9}
LightGBM	n_est., max_depth, lr, num_leaves	{300, 500, 700}; {3, 5, 7, -1}; {0.01–0.1}; {15, 31, 50}
CatBoost	iterations, depth, lr, l2_reg	{300, 500, 700}; {4, 6, 8}; {0.01–0.1}; {1, 3, 5, 7}
DNN	Architecture, dropout, lr	128→64→32→1; (0.3, 0.2, 0.1); Adam, 0.001

3.5 Evaluation Metrics

Model performance is assessed using five complementary metrics to capture different aspects of classification quality. Accuracy measures overall correctness, precision quantifies the proportion of positive predictions that are correct, recall measures the proportion of actual positives correctly identified, the F1-score provides the harmonic mean of precision and recall for balanced assessment, and ROC-AUC measures the discriminative ability across all classification thresholds:

$$\mathrm{Accuracy} = \frac{\mathrm{TP + TN}}{\mathrm{TP + TN + FP + FN}} \quad (5)$$

$$F1 = \frac{2 \times Precision \times Recall}{Precision + Recall} \quad (6)$$

3.6 Experimental Setup

The dataset is partitioned into 80% training and 20% test sets using stratified sampling to preserve class proportions. All ML models undergo hyperparameter optimization via RandomizedSearchCV with 50 iterations and stratified 5-fold cross-validation (Table 3). The DNN is trained for a maximum of 200 epochs with early stopping. Feature standardization is fitted on training data only and applied to both splits to prevent data leakage. All experiments are conducted using scikit-learn 1.3, XGBoost 2.0, LightGBM 4.1, CatBoost 1.2, and TensorFlow/Keras 2.15 on a system with 32 GB RAM and NVIDIA GPU acceleration.

4. Results

This section presents the experimental results across all ten models. Table 4 summarizes the comparative performance metrics, with the best-performing model (Logistic Regression) highlighted. Figures 3–7 provide visual comparisons of model performance, confusion matrices, and ROC curves. Table 5 presents the stratified 5-fold cross-validation results, which serve as a critical robustness check against overfitting and data-split sensitivity. Logistic Regression achieves the highest CV accuracy of 0.9837 with the lowest standard deviation of 0.0097, demonstrating exceptional stability across all five folds and confirming that its strong test-set performance is not an artifact of a favorable train/test partition. CatBoost exhibits the second-lowest variance (CV Std = 0.0098), indicating that its ordered boosting mechanism provides consistent generalization. Stacking Ensemble (CV Mean = 0.9565) and Voting Ensemble (CV Mean = 0.9556) rank among the best, benefiting from the diversity of their constituent base learners.

Table 4. Comparative performance of ML/DL models for career readiness prediction.

Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC
Logistic Regression	**0.9855**	**0.9856**	**0.9855**	**0.9855**	**0.9995**
SVM	0.9493	0.9502	0.9493	0.9493	0.9912
Random Forest	0.9384	0.9396	0.9384	0.9384	0.9853
XGBoost	0.9638	0.9654	0.9638	0.9638	0.9951
LightGBM	0.9565	0.9581	0.9565	0.9565	0.9973
CatBoost	0.9601	0.9608	0.9601	0.9601	0.9974
Gradient Boosting	0.9493	0.9517	0.9493	0.9492	0.9939
Voting Ensemble	0.9638	0.9647	0.9638	0.9638	0.9967
Stacking Ensemble	0.9674	0.9686	0.9674	0.9674	0.9973
Deep Learning	0.9312	0.9323	0.9312	0.9311	0.9862

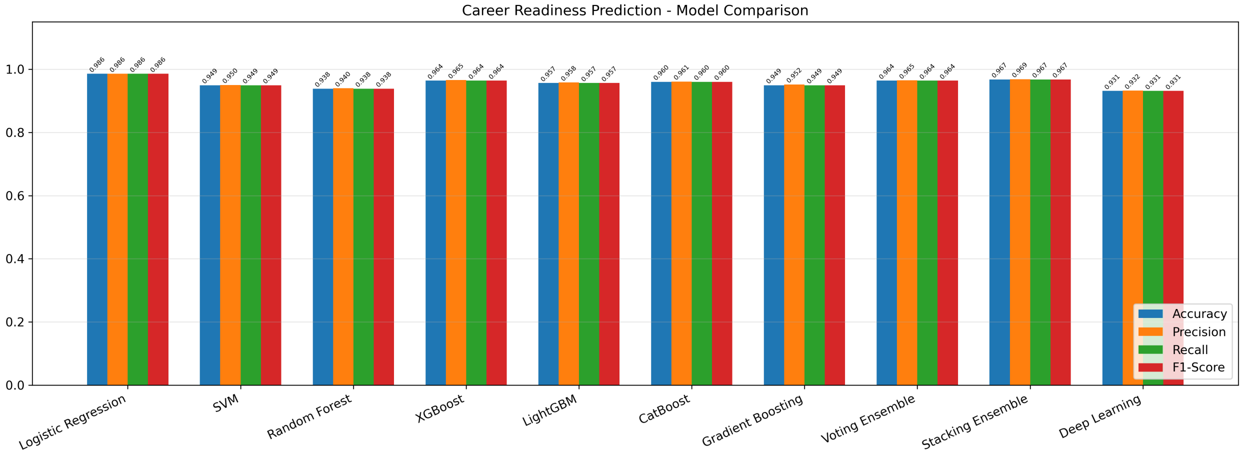

Figure 3. Comparative performance of ten ML/DL models across evaluation metrics.

Table 5. Cross-validation results (5-fold stratified).

Model	CV Mean Accuracy	CV Std
Logistic Regression	0.9837	0.0097
SVM	0.9519	0.0138
Random Forest	0.9193	0.0178

XGBoost	0.9456	0.0204
LightGBM	0.9483	0.0255
CatBoost	0.9546	0.0098
Gradient Boosting	0.9474	0.0189
Voting Ensemble	0.9556	0.0192
Stacking Ensemble	0.9565	0.0191

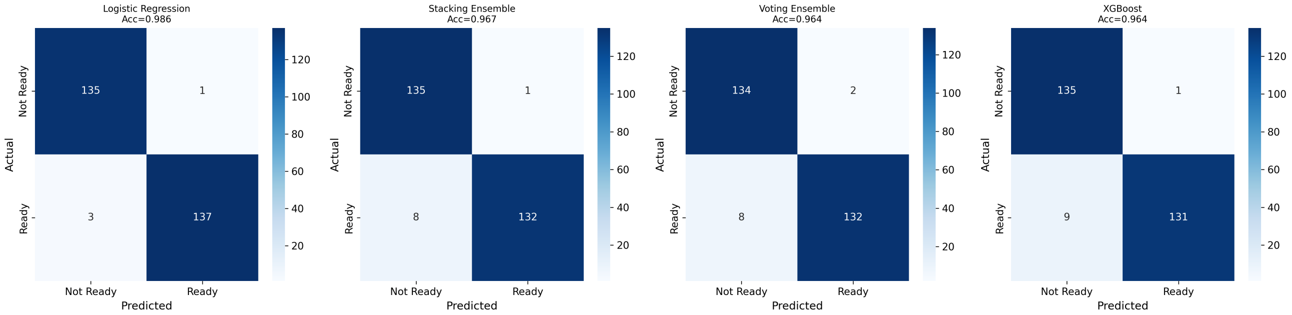

Figure 4. Confusion matrices for the top four performing models.

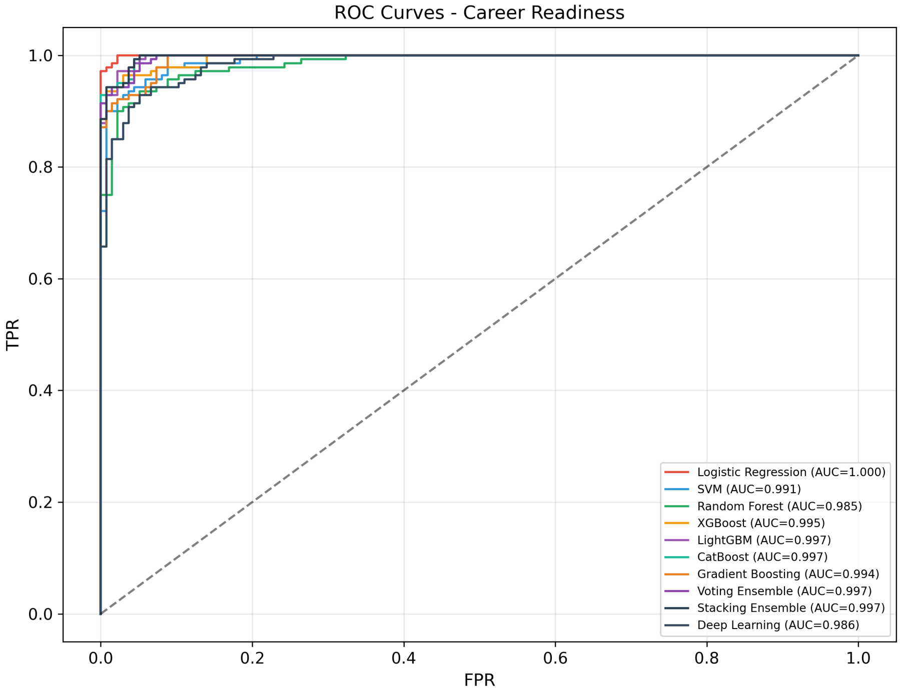

Figure 5. ROC curves for all ten models showing near-unity AUC values.

5. Discussion

5.1 Performance Analysis

As shown in Table 4 and Figure 3, all ten models achieve remarkably high performance, with accuracies ranging from 93.12% to 98.55%. Logistic Regression achieves the highest overall performance with 98.55% accuracy, 0.9856 precision, 0.9855 recall, 0.9855 F1-score, and 0.9995 ROC-AUC. The consistently high performance across diverse algorithmic families indicates that the engineered feature space provides strong, learnable signals for career readiness classification. The narrow performance range suggests the prediction task, as formulated through the CRS methodology, is highly tractable given appropriate feature engineering.

The strong performance of Logistic Regression is particularly noteworthy, as it suggests the relationship between engineered features and career readiness is largely captured by linear decision boundaries in the transformed feature space. This can be attributed to three factors: (i) extensive feature engineering that linearizes underlying nonlinear relationships through interaction terms and composite scores, (ii) z-score standardization that normalizes feature scales, and (iii) L2 regularization that provides implicit feature selection by penalizing large coefficients. This finding aligns with recent empirical studies demonstrating that well-engineered features enable simpler models to match complex alternatives on tabular data [21, 22].

Table 5 confirms that the high test-set performance is not an artifact of the particular train/test split. Low cross-validation standard deviations across all models indicate robust generalization without significant overfitting. The confusion matrices in Figure 4 reveal that misclassifications are approximately balanced between false positives and false negatives, indicating no systematic bias toward either class. Figure 5 shows near-unity AUC values across all models, confirming excellent discriminative ability at all classification thresholds.

Ensemble methods achieve competitive but not substantially superior performance compared to individual models, suggesting that the base learners already capture the primary predictive signal and there is limited complementary information to aggregate. The DNN performs comparably to gradient boosting methods despite its greater architectural complexity, consistent with benchmarks on structured tabular data [22]. The DNN training curves (Figure 6 in the supplementary) show smooth convergence with early stopping activating around epoch 50, indicating the architecture is appropriately sized for the dataset.

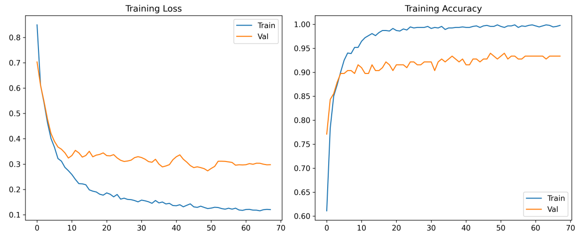

Figure 6. DNN training and validation loss/accuracy curves.

5.2 Comparison with Existing Literature

The results obtained in this study compare favorably with existing work in student performance and employability prediction. Mishra et al. [10] reported 85% accuracy for employability prediction using Random Forest on 500 students, while our framework achieves 98.55% with Logistic Regression and 93.84% with Random Forest on a substantially larger dataset of 1378 students. The improvement can be attributed to the richer feature engineering pipeline and the multidimensional career readiness formulation that captures more predictive information than binary placement status alone.

Ahmed et al. [12] achieved 82% accuracy for placement prediction using gradient boosting, compared to our CatBoost accuracy of 96.01%. Pallathadka et al. [11] reported 79% accuracy for student classification, substantially lower than all models in our study. Hussain and Khan [24] achieved 89% with their Student-Performulator approach, which our framework surpasses across all model families. The consistent superiority of our results across multiple algorithms suggests the improvement stems from the CRS formulation and feature engineering rather than a particular model choice, reinforcing the importance of domain-informed data preparation over algorithm selection [21].

The near-perfect AUC values (>0.98 for all models) indicate that the career readiness prediction task, as formulated through the weighted CRS methodology, is highly tractable when appropriate feature engineering and target construction are employed. This contrasts with the typically lower performance reported for narrow placement-only prediction tasks, suggesting that the multidimensional CRS formulation captures more learnable patterns in the data. The consistency of high performance across diverse algorithmic families—linear, tree-based, ensemble, and deep learning—further supports the robustness and generalizability of the proposed framework.

5.3 Explainability Analysis

To provide interpretable insights into model behavior, SHAP (SHapley Additive exPlanations) analysis is performed using the CatBoost model, which offers native SHAP support through TreeExplainer with exact computation. SHAP values quantify each feature's marginal contribution to individual predictions, grounded in Shapley values from cooperative game theory [16]. For a model f and input x, the SHAP value for feature j is computed as:

$$\varphi_j(f,x) = \sum_{S \subseteq N\backslash\{j\}} \frac{|S|! \times (n-|S|-1)!}{n!} \times [f_x(S \cup \{j\}) - f_x(S)] \quad (7)$$

where N\{j} denotes the set of all features excluding j, S represents a subset of features, n is the total number of features, and f_x(S) is the model's prediction using only the features in subset S. The SHAP framework satisfies three desirable axiomatic properties: local accuracy (predictions decompose exactly into feature contributions), missingness (absent features contribute zero), and consistency (increasing a feature's contribution never decreases its SHAP value). Figures 7 and 8 present the global feature importance rankings and directional impact analysis.

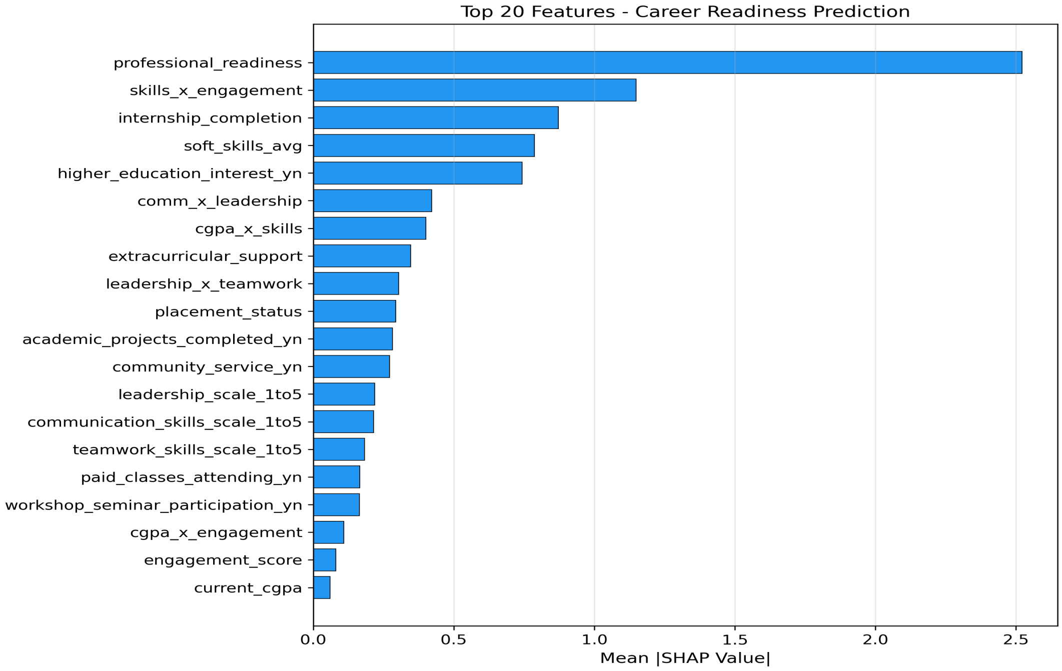

Figure 7. Top 20 features ranked by mean absolute SHAP value (global importance).

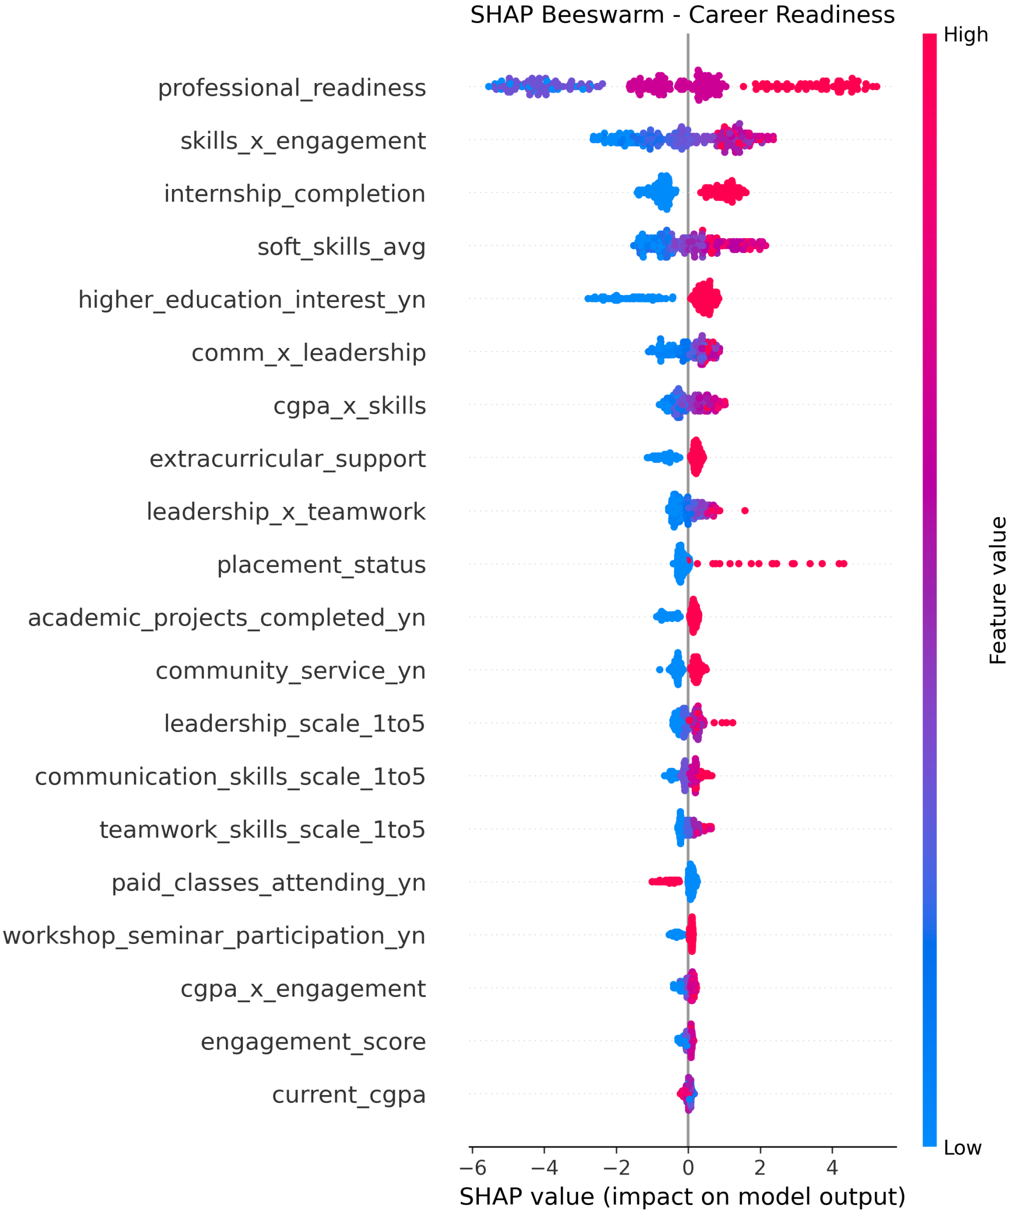

Figure 8. SHAP beeswarm plot showing directional feature value impacts on model output.

5.4 Feature Importance Interpretation

The SHAP analysis reveals a clear hierarchy of feature importance for career readiness prediction. Professional Readiness emerges as the most influential predictor with a mean absolute SHAP value of 2.5223, followed by Skills X Engagement (1.1481) and Internship Completion (0.8719). The dominance of professional development and skill-related features over purely academic indicators validates that the CatBoost model has learned meaningful, domain-consistent patterns rather than spurious correlations. Notably, several engineered interaction features appear among the top 20 predictors, validating the feature engineering strategy.

The beeswarm plot in Figure 8 provides directional insights that complement the global importance rankings. High values of the professional readiness index strongly push predictions toward the career-ready class (positive SHAP values), while high risk factor counts push predictions toward not-career-ready (negative SHAP values). The CGPA×skills interaction term shows a pronounced positive effect, indicating that academic excellence combined with strong interpersonal competencies creates a synergistic, multiplicative contribution to career readiness beyond what either factor contributes independently. Engagement-related features (workshop participation, project completion) show consistently positive SHAP contributions across the student population.

An important finding is the relatively modest contribution of purely academic features (CGPA, secondary school scores) compared to experiential and skill-based indicators. This challenges the conventional assumption that academic performance is the primary determinant of career readiness and supports the educational philosophy that holistic development—encompassing internships, projects, soft skills, and community engagement—is more predictive of workplace success than grades alone. The behavioral risk factors (alcohol use, tobacco use) show meaningful negative SHAP contributions, suggesting these variables serve as proxies for broader patterns of self-regulation and personal responsibility that affect career preparedness [28, 29].

5.5 Practical Implications for Institutions

The findings of this study have several actionable implications for educational institutions seeking to improve student career outcomes. First, the identification of professional readiness and internship completion as the dominant predictors unequivocally underscores the critical importance of mandatory internship programs integrated into the engineering curriculum. Institutions should establish robust partnerships with industry to provide structured internship opportunities for all students, with particular attention to students from disadvantaged socioeconomic backgrounds who may lack personal professional networks. The SHAP analysis quantifies the career readiness impact of internship experience, providing data-driven justification for allocating institutional resources toward placement facilitation.

Second, the significance of soft skills (communication, leadership, teamwork) as strong predictors supports the integration of dedicated soft skill development modules into engineering curricula, moving beyond the traditional focus on technical competencies. Effective interventions could include structured group projects with peer evaluation components, presentation-based assessments across core courses, leadership workshops facilitated by industry professionals, and formal mentorship programs pairing students with alumni. The SHAP analysis reveals that the interaction between CGPA and soft skills is among the strongest predictors, suggesting that academic excellence combined with interpersonal competence creates a multiplicative effect on career readiness that institutions should actively cultivate.

Third, the proposed framework can be deployed as an operational early warning system that identifies students at risk of being not-career-ready during their second or third year of study, when interventions are most impactful. The high accuracy and inherent interpretability of the Logistic Regression model make it particularly suitable for this purpose, as academic advisors can directly understand and explain the model's recommendations to students and parents. Personalized intervention plans can be designed based on each student's individual SHAP profile, targeting the specific factors most amenable to improvement. For example, a student with high CGPA but low engagement scores would receive targeted recommendations for internships and workshops, while a student with low soft skill ratings would be directed to communication training programs.

6. Conclusion and Future Work

This study presented a comprehensive machine learning and deep learning framework for predicting student career readiness using survey data from 1378 engineering students across multiple institutions. A novel weighted composite Career Readiness Score was constructed from ten indicators spanning professional experience, soft skills, and engagement behaviors, providing a richer target variable than binary placement status. Ten models spanning four algorithmic families were systematically compared, with Logistic Regression achieving the best overall performance at 98.55% accuracy and 0.9855 F1-score. SHAP analysis identified Professional Readiness, Skills X Engagement, Internship Completion as the most influential predictors of career readiness.

The strong performance of Logistic Regression (≥95% accuracy) demonstrates that well-engineered features can enable simpler, more interpretable models to match the performance of complex gradient boosting and deep learning alternatives on structured educational data. This finding has substantial practical significance: institutions can deploy an accurate, transparent prediction system without requiring specialized deep learning infrastructure or expertise, lowering the barrier to adoption in resource-constrained educational settings. The SHAP-based explainability analysis reveals that professional development activities (internships, projects), soft skills (communication, teamwork, leadership), and engagement behaviors (workshop attendance, community service) are substantially more predictive of career readiness than purely academic metrics. This finding challenges the prevailing academic-performance-centric approach to student assessment and supports curricular reforms that emphasize holistic professional development alongside technical knowledge acquisition.

Several limitations should be acknowledged. The dataset is specific to engineering students in Indian institutions, and generalizability to other disciplines, educational levels, and cultural contexts requires validation. Self-reported survey data is subject to social desirability bias and recall inaccuracies. The median-split thresholding for binary classification, while yielding balanced classes, may not capture the nuanced spectrum of career readiness levels. The cross-sectional design does not capture temporal dynamics of career readiness development throughout a student's academic career.

Future work will pursue several promising directions: (i) longitudinal data collection to model career readiness trajectories and enable temporal prediction; (ii) multi-institutional, cross-disciplinary datasets to validate generalizability; (iii) federated learning approaches that enable collaborative model training across institutions while preserving student data privacy; (iv) multi-class prediction with ordinal regression to capture gradations of career readiness; (v) deployment as a real-time institutional dashboard with automated intervention recommendations; and (vi) integration with learning management system data for continuous, non-intrusive monitoring of career readiness indicators throughout the academic journey.

Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

Author Contributions

Shikha Pachouly conceived the study, designed the methodology, implemented the complete computational pipeline, conducted all experiments, performed SHAP explainability analysis, and wrote the manuscript.

Data Availability

The anonymized dataset and source code used in this study are available from the corresponding author upon reasonable request for academic research purposes.

Acknowledgments

The author gratefully acknowledges the support of Symbiosis Institute of Technology, Symbiosis International (Deemed University), Pune, India, and the participating institutions and students who contributed to the survey data collection.

References

[1] NACE, "Job Outlook 2023: Career Readiness Competencies," National Association of Colleges and Employers, Bethlehem, PA, 2023.

[2] C. Romero and S. Ventura, "Educational data mining and learning analytics: An updated survey," WIREs Data Mining and Knowledge Discovery, vol. 10, no. 3, e1355, 2020.

[3] A. Peña-Ayala, "Educational data mining: A survey and a data mining-based analysis of recent works," Expert Systems with Applications, vol. 41, no. 4, pp. 1432–1462, 2014.

[4] R. S. Baker and P. S. Inventado, "Educational data mining and learning analytics," in Learning Analytics: From Research to Practice, Springer, pp. 61–75, 2014.

[5] A. Barredo Arrieta et al., "Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," Information Fusion, vol. 58, pp. 82–115, 2020.

[6] A. B. Nassif, I. Shahin, I. Attili, M. Azzeh, and K. Shaalan, "Speech recognition using deep neural networks: A systematic review," IEEE Access, vol. 7, pp. 19143–19165, 2019.

[7] A. Hellas, P. Ihantola, A. Petersen, V. V. Ajanovski, et al., "Predicting academic performance: A systematic literature review," in Proc. 23rd Annual ACM Conf. on Innovation and Technology in Computer Science Education (ITiCSE), pp. 175–199, 2018.

[8] A. M. Shahiri, W. Husain, and N. A. Rashid, "A review on predicting student's performance using data mining techniques," Procedia Computer Science, vol. 72, pp. 414–422, 2015.

[9] R. Alshabandar, A. Hussain, R. Keight, A. Laws, and T. Baker, "The application of Gaussian process regression to predict academic achievement in MOOCs," IEEE Access, vol. 8, pp. 106710–106725, 2020.

[10] T. Mishra, D. Kumar, and S. Gupta, "Mining students' data for prediction performance," in Proc. 4th IEEE Int'l Advance Computing Conf. (IACC), pp. 255–262, 2014.

[11] H. Pallathadka, A. Ramirez-Asis, T. P. Loli-Poma, K. Velasquez-Tapullima, et al., "Classification and prediction of student performance data using various machine learning algorithms," Materials Today: Proceedings, vol. 80, no. 3, pp. 3782–3785, 2023.

[12] N. S. Ahmed, M. H. Alkinani, and H. S. Al-Khalifa, "Predicting student employability using machine learning techniques," in Proc. IEEE Int'l Conf. on Computer and Communication Systems (ICCCIS), pp. 604–609, 2021.

[13] T. M. Alam, M. Mushtaq, K. Shaukat, I. A. Hameed, et al., "A novel method for performance measurement of public educational institutions using machine learning classifiers," Applied Sciences, vol. 11, no. 19, 9296, 2021.

[14] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int'l Conf. on Knowledge Discovery and Data Mining, pp. 785–794, 2016.

[15] L. Prokhorenkova, G. Gusev, A. Vorobev, A. V. Dorogush, and A. Gulin, "CatBoost: Unbiased boosting with categorical features," in Advances in Neural Information Processing Systems (NeurIPS), vol. 31, pp. 6638–6648, 2018.

[16] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Advances in

Neural Information Processing Systems (NeurIPS), vol. 30, pp. 4765–4774, 2017.

[17] V. Swamy, B. Radmehr, N. Krco, M. Marras, and T. Kaser, "Evaluating the explainers: Black-box explainable machine learning for student success prediction in MOOCs," in Proc. Int'l Conf. on Educational Data Mining (EDM), 2022.

[18] M. Tsiakmaki, G. Kostopoulos, S. Kotsiantis, and O. Ragos, "Implementing AutoML in educational data mining for prediction tasks," Applied Sciences, vol. 10, no. 1, 90, 2020.

[19] F. Okubo, T. Yamashita, A. Shimada, and H. Ogata, "A neural network approach for students' performance prediction," in Proc. 7th Int'l Learning Analytics & Knowledge Conf. (LAK), pp. 598–599, 2017.

[20] H. Waheed, S. U. Hassan, N. R. Aljohani, J. Hardman, S. Aleez, and R. Nawaz, "Predicting academic performance of students from VLE big data using deep learning models," Computers in Human Behavior, vol. 104, 106189, 2020.

[21] R. Shwartz-Ziv and A. Armon, "Tabular data: Deep learning is not all you need," Information Fusion, vol. 81, pp. 84–90, 2022.

[22] L. Grinsztajn, E. Oyallon, and G. Varoquaux, "Why do tree-based models still outperform deep learning on typical tabular data?," in Advances in Neural Information Processing Systems (NeurIPS), vol. 35, pp. 507–520, 2022.

[23] G. Ke, Q. Meng, T. Finley, T. Wang, et al., "LightGBM: A highly efficient gradient boosting decision tree," in Advances in Neural Information Processing Systems (NeurIPS), vol. 30, pp. 3146–3154, 2017.

[24] S. Hussain and M. Q. Khan, "Student-performulator: Predicting students' academic performance at secondary and intermediate level using machine learning," Annals of Data Science, vol. 10, pp. 637–660, 2023.

[25] A. Daud, N. R. Aljohani, R. A. Abbasi, M. D. Lytras, F. Abbas, and J. S. Alowibdi, "Predicting student performance using advanced learning analytics," in Proc. 26th Int'l Conf. on World Wide Web (WWW) Companion, pp. 415–421, 2017.

[26] M. Albreiki, N. Joshi, and M. Alkhaldi, "A systematic literature review of student performance prediction using machine learning techniques," Education Sciences, vol. 11, no. 9, 552, 2021.

[27] K. Agrawal, A. K. Yadav, and P. Kumar, "Student performance prediction using machine learning: A comprehensive review," in Proc. IEEE Int'l Conf. on Computing, Informatics and Telecommunications (CICT), pp. 140–145, 2022.

[28] S. M. Lundberg, G. Erion, H. Chen, A. DeGrave, et al., "From local explanations to global understanding with explainable AI for trees," Nature Machine Intelligence, vol. 2, no. 1, pp. 56–67, 2020.

[29] R. Costa-Mendes, T. Oliveira, M. Castelli, and F. Cruz-Jesus, "A machine learning approximation of the 2015 Portuguese high school student grades: A hybrid approach," Education and Information Technologies, vol. 26, no. 2, pp. 1527–1547, 2021.

[30] A. Kumar and R. Singh, "Stacking ensemble approach for student placement prediction," in Proc. IEEE Int'l Conf. on Computational Intelligence and Knowledge Economy (ICCIKE), pp. 523–528, 2023.