

Article

Comprehensive Feature Analysis and Evaluation on the Student Performance Based on Machine Learning

Zhifeng Zhang ¹, Xiaoyun Qin ², Yangyang Chu ¹, Junxia Ma ¹  and Bo Wang ^{3,*} ¹ Software Engineering College, Zhengzhou University of Light Industry, Zhengzhou 450000, China² Department of Material and Chemical Engineering, Zhengzhou University of Light Industry, Zhengzhou 450000, China³ School of Computer Science and Technology, Henan Institute of Technology, Xinxiang 453003, China

* Correspondence: wangb@hait.edu.cn

Abstract

The cultivation of high-quality talents relies on the synergistic interaction of various educational stakeholders, including schools, families, and society, within the educational system. With the rapid advancement of artificial intelligence (AI) technology, new opportunities have emerged for constructing and optimizing collaborative education mechanisms. Based on a feature-rich and large-scale real dataset, this paper conducts a case study to explore novel approaches for leveraging AI to empower a cooperative education system. Specifically, correlation and association analysis methods from traditional statistics are first employed to quantify pairwise feature relationships, providing a basis for identifying key factors influencing student development. Subsequently, principal component analysis (PCA) is applied to extract dominant components from the dataset, assess the intrinsic information carried by each feature, and uncover latent relationships among features. Finally, leveraging the multi-source and heterogeneous nature of the cooperative education system, a novel multi-branch neural network model (MBDNN) is proposed to achieve accurate prediction of student academic performance. This study can provide reference and methodological support for effectiveness evaluation and decision-making within the cooperative education system.

Keywords: cooperative education; talent cultivation; artificial intelligence; deep neural network; association analysis

1. Introduction

The cultivation of high-quality talents is a complex systemic project, whose effectiveness highly depends on the synergistic interactions among multiple educational entities including students, schools, families, and society [1–3]. This school–family–society collaborative education system (SFSCES) framework recognizes that student development emerges from the interplay of academic environments, family backgrounds, and socioeconomic contexts—each contributing unique and often nonlinear influences. However, traditional educational analysis methods are often based on isolated and linear evaluation models, which struggle to effectively capture the complex nonlinear interactions among multiple factors, and thus fail to provide comprehensive and precise decision-making support for the systematic optimization of collaborative education strategies.

In recent years, the continuous development of artificial intelligence (AI) technology, particularly breakthroughs in data mining and deep learning [4], has provided new research paradigms and technical pathways for in-depth analysis of complex educational



Academic Editor: Juan Francisco De Paz Santana

Received: 1 June 2026

Revised: 23 June 2026

Accepted: 26 June 2026

Published: 2 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

phenomena and the construction of intelligent educational support systems [5–7]. However, a critical challenge persists: many state-of-the-art deep learning models in educational data mining (EDM) prioritize predictive accuracy at the expense of interpretability, functioning as “black boxes” that obscure the reasoning behind their predictions [8]. Recent studies have attempted to address this through hybrid models combining CNNs with attention mechanisms, deep ensemble stacking, or graph neural networks, yet these approaches often remain opaque to educators and administrators who need to understand why a particular student receives a specific prediction [9,10]. This interpretability gap is particularly consequential in educational settings, where predictions inform high-stakes decisions (early warning systems, resource allocation, and personalized interventions) and where trust and transparency are prerequisites for adoption [11].

This study aims to use AI algorithms to systematically analyze the multidimensional SFSCES environment and key factors affecting student academic performance, thereby providing a scientific basis for precise intervention, resource allocation, and strategic decision-making within the SFSCES [12,13]. The research consists of three main parts. First, traditional statistical methods are used to quantitatively analyze pairwise feature relationships and preliminarily identify key influencing factors, including Pearson and Spearman correlation coefficients [14] as well as Cramer’s V, Tschuprow’s T, and Pearson’s contingency coefficients [15]. Second, principal component analysis (PCA) is employed to reduce feature dimensionality and enable reconstruction, uncovering latent covariance patterns among features and assessing the global importance of each feature. Third, building on the above analysis, a novel multi-branch deep neural network (MBDNN) model designed for multi-source collaborative data is proposed to accurately predict student academic performance.

The proposed MBDNN, while structurally lightweight with only two hidden layers, embodies a problem-driven design philosophy that distinguishes it from existing deep learning approaches in EDM. Unlike complex black-box models that indiscriminately mix all features, the MBDNN’s three branches are explicitly designed to correspond to the school, family, and societal subsystems of the SFSCES, a mapping that is both theoretically grounded and empirically validated by PCA (the first three principal components explaining over 85% of the variance). This design prioritizes interpretability as a first-class principle, enabling stakeholders to trace which subsystem contributed most to a particular prediction. While recent studies have proposed multi-branch CNNs or hybrid architectures for student prediction, their branch designs are typically data-driven rather than theory-driven; the MBDNN represents a deliberate alternative that embeds domain knowledge directly into the network architecture, bridging the gap between predictive performance and educational explainability.

Based on the real student performance dataset provided by the Polytechnic Institute of Portalegre (comprising 36 features and 4424 samples) [16,17], this study yields the following key findings through systematic analysis. First, there are significant correlations and information redundancy among multiple student features, with prior academic performance, family background (such as parental education and occupation), and macro socioeconomic environment (e.g., GDP) identified as key factors influencing student development. Second, PCA results show that the first three principal components collectively explain over 85% of the total variance, further confirming strong correlations among features. Third, experiments demonstrate that the proposed MBDNN model significantly outperforms various machine learning models in predictive accuracy (with an improvement ranging from 5.23% to 15.3%), and its multi-branch mechanism effectively captures the synergistic effects of factors from schools, families, and society. These findings not only validate the argument that student development is influenced by multi-party collaboration from a data perspec-

tive, but also provide quantifiable evidence and advanced analytical tools for educational administrators to evaluate effectiveness and formulate precise educational strategies.

The rest of the paper is organized as follows: Section 2 details the dataset used and preprocessing methods, and elaborates on the feature correlation analysis, PCA-based importance evaluation and the design of the MBDNN model. Section 3 presents experimental results and provides analytical discussion. Section 4 discusses the related works. Finally, Section 5 summarizes the research findings and outlines directions for future work.

2. Materials and Methods

Based on a real-world dataset of student academic status, this study begins by employing variable association analysis and feature compression methods to explore intrinsic relationships among features, with the aim of comprehensively identifying key factors affecting student performance and providing empirical evidence for the collaborative education mechanism involving schools, families, and society. Subsequently, deep learning techniques are adopted to construct a high-accuracy prediction model for student performance, thereby offering reliable support for decision evaluation and program design within the cooperative education system. The specific technical roadmap is illustrated in Figure 1.




Figure 1. Technology roadmap for AI-empowered cooperative education with a real dataset case.

For the given dataset, each feature is first subjected to normalization, scaling the value range of all features into a unified space to prevent bias caused by differences in numerical scales during association analysis and relational modeling. Considering the uneven distribution of attribute data (as shown in Figure 2) and the presence of certain outliers, robust normalization is applied instead of standard normalization to mitigate the impact of outliers on the normalization process. Specifically, for each sample value x in any feature, it is transformed into x_{scaled} according to Equation (1), where $x_{1/2}$ represents the median of the feature, and $x_{1/4}$ and $x_{3/4}$ denote the 25th and 75th percentile values of the feature, respectively.

$$x_{scaled} = \frac{x - x_{1/2}}{x_{3/4} - x_{1/4}} \quad (1)$$

Next, variable association analysis techniques are used to conduct a comprehensive and systematic evaluation of the features in the dataset, delving into the inherent associations between every two features. Building on this, PCA is applied for dimensionality reduction to eliminate features with high multicollinearity and low contribution, thereby accurately identifying key factors influencing the target variable (student performance) and providing clear and reliable data support for the formulation of collaborative education policies.

Finally, based on the above results and addressing the multi-class nature of student performance, a multi-branch deep neural network model is constructed to capture the complex nonlinear mapping relationships between multi-source student features and the target variable, offering a reliable tool for predicting academic status. This model helps to identify critical factors influencing student performance, and uncover the essential patterns related to the collaborative education mechanism behind the complex data.




Figure 2. The distribution of variables in the student performance dataset.

2.1. Dataset

In selecting the dataset, this study adhered to the following principles:

- **Authenticity of Data:** The data were sourced from real collected records rather than being generated through simulation or synthesis, ensuring the research findings are grounded in reality and possess generalizability.
- **Sufficient Data Scale:** The dataset is sufficiently large to cover diverse student cases, thereby enhancing the generalizability of both the analytical results and models.
- **Broad Feature Coverage:** The data include rich attribute information encompassing key factors that influence student performance.
- **Diverse Attribute Sources:** The attributes cover educational environment data provided by multiple sources including schools, families, and society, as well as multifaceted characteristics of the students themselves, offering a comprehensive reflection of students' development backgrounds and individual profiles.

Accordingly, this study utilizes a dataset from the Polytechnic Institute of Portalegre [16,17], spanning academic performance records of students from the 2005/06 to the 2018/19 academic years at the end of their first and second semesters. The dataset comprises 4424 samples, covering 36 feature variables and one target variable, as shown in Table 1. The features primarily consist of information known at the time of student enrollment, including academic background, demographic details, and socioeconomic factors. The target variable is categorical, comprising three classes: dropout, enrolled, and graduate.

Table 1. Variables of the student performance dataset.

Notation	Variable	Description
F1	Marital Status	The student's marital status, coded numerically (e.g., 1—single, 2—married).
F2	Application mode	The method of application used by the student, coded numerically for various application phases and special contingents.
F3	Application order	The student's application preference, ranging from 0 (first choice) to 9 (last choice).
F4	Course	The specific undergraduate course the student is enrolled in, represented by numerical codes for various fields such as Agronomy, Design, Education, Nursing, Journalism, Management, Social Services, and Technology.
F5	Daytime/evening attendance	Indicates if the student attends classes during the day (1) or in the evening (0).
F6	Previous qualification	The qualification obtained by the student prior to higher education enrollment, coded numerically for different levels (e.g., secondary education, bachelor's degree).
F7	Previous qualification (grade)	The grade achieved in the previous qualification, ranging from 0 to 200.
F8	Nationality	The student's nationality, represented by numerical codes (e.g., 1—Portuguese, 41—Brazilian).
F9	Mother's qualification	The educational qualification of the student's mother, coded numerically.
F10	Father's qualification	The educational qualification of the student's father, coded numerically.
F11	Mother's occupation	The occupation of the student's mother, coded numerically for various professional categories.
F12	Father's occupation	The occupation of the student's father, coded numerically for various professional categories.
F13	Admission grade	The student's admission grade, ranging from 0 to 200.
F14	Displaced	Indicates if the student is a displaced person (1—yes, 0—no).
F15	Educational special needs	Indicates if the student has any special educational needs (1—yes, 0—no).
F16	Debtor	Indicates if the student is a debtor (1—yes, 0—no).
F17	Tuition fees up to date	Indicates if the student's tuition fees are up to date (1—yes, 0—no).
F18	Gender	The student's gender (1—male, 0—female).
F19	Scholarship holder	Indicates if the student is a scholarship holder (1—yes, 0—no).
F20	Age at enrollment	The student's age at the time of enrollment.
F21	International	Indicates if the student is an international student (1—yes, 0—no).
F22	Curricular units 1st sem (credited)	The number of curricular units credited in the first semester.
F23	Curricular units 1st sem (enrolled)	The number of curricular units the student enrolled in during the first semester.
F24	Curricular units 1st sem (evaluations)	The number of evaluations for curricular units in the first semester.

Table 1. Cont.

Notation	Variable	Description
F25	Curricular units 1st sem (approved)	The number of curricular units approved in the first semester.
F26	Curricular units 1st sem (grade)	The average grade in the first semester, between 0 and 20.
F27	Curricular units 1st sem (without evaluations)	The number of curricular units without evaluations in the first semester.
F28	Curricular units 2nd sem (credited)	The number of curricular units credited in the second semester.
F29	Curricular units 2nd sem (enrolled)	The number of curricular units the student enrolled in during the second semester.
F30	Curricular units 2nd sem (evaluations)	The number of evaluations for curricular units in the second semester.
F31	Curricular units 2nd sem (approved)	The number of curricular units approved in the second semester.
F32	Curricular units 2nd sem (grade)	The average grade in the second semester, between 0 and 20.
F33	Curricular units 2nd sem (without evaluations)	The number of curricular units without evaluations in the second semester.
F34	Unemployment rate	The unemployment rate (%).
F35	Inflation rate	The inflation rate (%).
F36	GDP	Gross Domestic Product (GDP).
T	Target	The classification target, indicating if the student is 'dropout', 'enrolled', or 'graduate' at the end of the course.

2.2. Feature Analysis Tools

In real-world collaborative education scenarios, student performance is influenced by multiple factors, encompassing individual characteristics as well as the environment and resources provided by the family, school, and society. These factors exhibit multidimensional and highly nonlinear complex relationships. Therefore, relying solely on a single type of correlation analysis method is insufficient to fully capture such intricate interdependencies. Unlike many existing studies that primarily use Pearson correlation coefficients, this paper comprehensively employs multiple correlation quantification methods to systematically and multidimensionally reveal the relationships among features within SFSCES, thereby providing a solid data foundation and decision-making basis for optimizing collaborative education strategies.

Specifically, this study first adopts two types of traditional statistical correlation coefficients (Pearson and Spearman rank correlation coefficients) along with three association measures for categorical variables (Cramer's V, Tschuprow's T, and Pearson's contingency coefficient). These methods quantitatively evaluate pairwise feature associations from different perspectives, including linear and nonlinear, parametric and nonparametric approaches. By integrating results from multiple methods, a comprehensive correlation metric is constructed to enhance the robustness and comprehensiveness of association identification, laying a reliable foundation for subsequent feature selection and predictive modeling.

Furthermore, principal component analysis (PCA) is introduced to perform linear transformation and dimensionality reduction on the original features, extracting low-dimensional feature representations that retain maximum explanatory variance. The resulting principal components not only preserve the essential discriminatory information from the original data but also provide critical insights for identifying key influencing factors and improving model interpretability.

The following subsections detail the two technical schemes in detail, pairwise feature association analysis and PCA-based feature analysis.

2.2.1. Feature Association Analysis

First, this study employs the Pearson correlation coefficient to measure the linear correlation between two features. Given observed samples of two features, $A = [a_1, a_2, \dots, a_n]$

and $B = [b_1, b_2, \dots, b_n]$, the Pearson correlation coefficient ρ_{AB}^P is defined as Equation (2). $\bar{A}$ represents the average of $a_i, i = 1, \dots, n$, i.e., $\bar{A} = \sum_{i=1}^n a_i / n$.

$$\rho_{AB}^P = \frac{\sum_{i=1}^n (a_i - \bar{A})(b_i - \bar{B})}{\sqrt{\sum_{i=1}^n (a_i - \bar{A})^2} \sqrt{\sum_{i=1}^n (b_i - \bar{B})^2}} \quad (2)$$

Then, the Spearman rank correlation coefficient is used to assess the strength of the rank correlation between two features, which reflects the degree of association describable by a monotonic function. The Spearman rank correlation coefficient (ρ_{AB}^S) between two features is calculated as shown in Equation (3). $R(a_i)$ represents the rank of a_i in A . $\bar{R(A)}$ is the average of $R(a_i), i = 1, \dots, n$.

$$\rho_{AB}^S = \frac{\sum_{i=1}^n (R(a_i) - \bar{R(A)})(R(b_i) - \bar{R(B)})}{\sqrt{\sum_{i=1}^n (R(a_i) - \bar{R(A)})^2} \sqrt{\sum_{i=1}^n (R(b_i) - \bar{R(B)})^2}} \quad (3)$$

The above correlation coefficients are computed based on numerical feature samples and are suitable for analyzing relationships between continuous numerical features. To effectively measure the association strength between discrete features, this study further introduces distribution-based statistical measures of variable association, specifically including Pearson's, Cramer's V, and Tschuprow's T contingency coefficients.

Given observed samples of two features $(A, B) = \{(a_1, b_1), (a_2, b_2), \dots, (a_n, b_n)\}$, we use n_{ab} to represent the frequency that (a, b) is observed. $n_{a\cdot}$ and $n_{\cdot b}$ represent the frequency that a and b are observed in A and B , respectively. Then, Pearson's, Cramer's V, and Tschuprow's T contingency coefficients can be calculated by Equations (4), (5) and (6), respectively, where χ^2 is the Pearson's chi-squared value, which is calculated by Equation (7). r and k represent the number of categories for the two features, respectively.

$$P_{AB} = \sqrt{\frac{\chi^2/n}{1 + \chi^2/n}} \quad (4)$$

$$V_{AB} = \sqrt{\frac{\chi^2/n}{\min(r-1, k-1)}} \quad (5)$$

$$T_{AB} = \sqrt{\frac{\chi^2/n}{\sqrt{(r-1)(k-1)}}} \quad (6)$$

$$\chi^2 = \sum_{(a,b) \in (A,B)} \frac{n_{ab} - n_{a\cdot}n_{\cdot b}/n}{n_{a\cdot}n_{\cdot b}/n} \quad (7)$$

For any two features, if the absolute value of any of the aforementioned five quantitative measures of association strength is close to 1, it indicates a strong association between the two features, further implying a high degree of redundancy in the information they provide. In such cases, when performing key factor analysis and building relationship models, it may be advisable to remove one of the features to reduce model complexity and help mitigate the risk of overfitting. That is to say, if any one of the five association measures indicates a strong relationship between two features, then those two features should be regarded as strongly associated. Therefore, in the subsequent feature association analysis, this study selects the maximum value among the five association measures as a comprehensive quantitative basis for the interdependency between features. That is to

say, given two features A and B , their comprehensive relational degree is quantified by Equation (8) in this paper.

$$\rho_{AB} = \max\{|\rho_{AB}^P|, |\rho_{AB}^S|, P_{AB}, V_{AB}, T_{AB}\} \quad (8)$$

On the other hand, if a feature exhibits a strong correlation with the target variable, its potential contribution to the target task is generally greater, making it a strong candidate in feature selection. However, in practical applications, it is essential to carefully distinguish between correlation and causality, examine redundancy among features (multicollinearity), and employ appropriate correlation measurement methods to ensure that true predictive relationships are captured, thereby enabling more comprehensive and robust decision-making.

2.2.2. Principal Component Analysis-Based Feature Analysis

The aforementioned traditional statistical association analysis methods often rely on strong distributional assumptions and can only handle pairwise feature relationships, making it difficult to effectively capture complex interactions among multiple features. To address this, this study introduces PCA, which performs an orthogonal transformation on the original features, converting them into a set of linearly uncorrelated principal components. This reduces the data dimensionality while retaining most of the original information, thereby effectively identifying structures and patterns implicitly shared across multiple features.

PCA is a dimensionality reduction algorithm widely used in data feature analysis. Its core idea is to map the original n -dimensional features into a k -dimensional ($k < n$) subspace via linear transformation. This subspace consists of a set of mutually orthogonal new features (the principal components) ensuring that the projection of the data in this subspace retains the maximum variance information from the original data. By preserving the top k principal components with the highest variance contributions, we can focus the analysis of high-dimensional data on a few representative features, thereby improving the efficiency and interpretability of modeling.

Given a vector of n features $\mathbf{X}$, PCA first computes its covariance matrix and performs eigenvalue decomposition. Then, the top k largest eigenvalues ($\Lambda = [\lambda_1, \lambda_2, \dots, \lambda_k]$) and their corresponding eigenvectors are selected to form the linear transformation matrix $\mathbf{T} \in \mathbb{R}^{k \times n}$. Finally, through the transformation $\mathbf{Y} = \mathbf{TX}$, the original n -dimensional features are mapped to k new features (principal components) $\mathbf{Y} = [Y_1, Y_2, \dots, Y_k]$. This transformation aims to preserve as much variation from the original data as possible, minimizing information loss during dimensionality reduction.

For the principal components obtained via PCA transformation, the corresponding eigenvalues reflect the amount of variance (represent the amount of information to a certain extent) contained in each principal component. Therefore, the proportion of original information carried by each principal component can be measured by the percentage of its corresponding eigenvalue relative to the sum of all eigenvalues (variance contribution rate). That is, the proportion of original information explained by principal component Y_j can be calculated by $\frac{\lambda_j}{\sum_{i=1}^n \lambda_i}$.

During the linear transformation process, the absolute value of the weight of each original feature in the principal components (i.e., the elements in the transformation matrix) reflects the degree of influence of that original feature on the formation of the principal component. To comprehensively evaluate the importance of an original feature in the reduced-dimensional space, a weighted sum of the absolute values of its weights across all principal components, weighted by the corresponding eigenvalues, can be computed to estimate its overall contribution to the feature representation (importance measure).

Therefore, this study uses the method defined in Equation (9) to evaluate the importance of feature X_i , where t_{ji} is the element at row j and column i of T . It should be noted that, since the focus of this study is on analyzing the relative importance among features rather than their absolute values, the importance measures were not normalized during computation. This strategy helps to more clearly identify key influencing factors and their relative contribution differences in subsequent analyses.

$$I_i = \sum_{j=1}^k (\lambda_j t_{ji}) \quad (9)$$

The importance metric defined in Equation (9) is grounded in established PCA theory. In PCA, each eigenvalue λ_j measures the amount of variance explained by principal component j , while the loadings t_{ji} quantify the contribution (correlation strength) of each original variable to that principal component. A larger absolute loading indicates that the feature contributes more significantly to that PC. Therefore, aggregating loadings across principal components weighted by their respective eigenvalues naturally yields a composite measure of each feature's overall importance to the variance structure of the data.

This weighted aggregation approach is well-documented in the literature. Jolliffe [18] established the theoretical foundation for using principal component loadings in identifying subsets of original variables that are strongly related to the first few PCs. Cadima and Jolliffe [19] further demonstrated that loading-based feature assessment is a principled approach for variable importance evaluation. More recent works have explicitly adopted this weighted aggregation strategy. For example, the SurvMarker R package [20] quantifies feature importance by aggregating loadings across all survival-associated PCs, weighted by the variance explained by each PC. Similarly, unsupervised feature selection methods using weighted principal components combine the first k PCs into a weighted composite where each loading reflects the contribution of each individual feature [21]. Thus, Equation (9) follows established practice in the PCA literature for aggregating multi-component information into a single interpretable importance score.

2.2.3. Limitations of Correlation- and PCA-Based Analyses

It is essential to recognize that both pairwise association measures (Pearson, Spearman, Cramer's V, etc.) and PCA are descriptive statistical tools that capture patterns of dependency and variance, but they do not establish causal relationships. A strong correlation between two features may arise from confounding factors, reverse causality, or measurement artifacts, rather than from a direct causal effect. Similarly, the feature importance scores derived from PCA (Equation (9)) reflect statistical influence on the variance explained by principal components, not causal impact; a high score does not imply that intervening on that feature would necessarily improve student outcomes.

In this study, these techniques are employed exclusively for exploratory and predictive purposes: feature screening, redundancy detection, dimensionality reduction, and revealing latent data structures to guide model design (e.g., the three-branch architecture of MBDNN aligns with the three dominant principal components). They are not intended to support causal claims. Readers are therefore cautioned against overinterpreting the findings as direct evidence for policy interventions. Translating these associations into actionable educational policies requires additional causal methodologies, such as instrumental variable analysis, regression discontinuity, or randomized experiments, which we plan to explore in future work.

2.3. Multi-Branch Deep Neural Network-Based Prediction Model

To quantify the effectiveness of collaborative education policies and provide data-driven decision support for formulating new educational strategies, this study employs the deep learning technique to construct a student performance prediction model. The model aims to accurately predict student development outcomes based on multidimensional inputs, including individual characteristics, environmental factors, and available resources, thereby offering a scientific basis for optimizing collaborative education strategies.

This paper proposes a multi-branch deep neural network (MBDNN, as shown in Figure 3) designed for this three-class prediction task. The model consists of an input layer, a dual hidden-layer structure with three parallel branches, and an output layer. The input layer receives a 36-dimensional feature vector (after preprocessing) plus one bias term. The first hidden layer contains 10 neurons with ReLU activation, organized into three parallel branches according to feature source and semantic meaning, as detailed in Table 2. The outputs of the three branches are concatenated and fed into a second hidden layer with three neurons (ReLU activation), which then connects to a three-neuron output layer with Softmax activation to produce class probabilities.




Figure 3. The multi-branch deep neural network.

Table 2. Feature-to-branch assignment in the MBDNN architecture.

Branch	Feature Group	Specific Features	Rationale
Branch 1	School/Academic Performance	F22–F26, F28–F32 (curricular units: credited, enrolled, evaluations, approved, grade, for both semesters)	Captures academic engagement, progress, and achievement within the educational institution; features share strong logical and statistical dependencies (Section 3.1, Group 2).
Branch 2	Family Background	F9–F12 (mother’s qualification, father’s qualification, mother’s occupation, father’s occupation)	Represents family socioeconomic and cultural capital; features exhibit strong inter-correlations (Section 3.1, Group 1) and directly reflect family system influences.
Branch 3	Socioeconomic and Other Factors	F34–F36 (unemployment rate, inflation rate, GDP), plus F1–F8, F13–F21, F27, F33	Includes macroeconomic indicators (reflecting societal conditions) and other individual characteristics not belonging to the first two categories.

The MBDNN's multi-branch design, while structurally simple, embodies a principled innovation that lies not in introducing novel computational primitives, but in the purposeful, theoretically grounded, and data-validated design of the architecture specifically tailored to the collaborative education problem. This innovation manifests in four key aspects. First, the branch design is domain-specific and theory-driven: unlike generic multi-branch networks that use branches for different kernel sizes (e.g., Inception) or different modalities (e.g., vision + text), MBDNN's branches are explicitly defined based on the SFSCES framework, with each branch processing a semantically coherent group of features that correspond to one of the three educational subsystems. This embeds domain knowledge directly into the network architecture, an approach rarely seen in existing educational data mining studies. Second, the number of branches and the dimensionality of the fusion layer are empirically validated by PCA: the results presented in Section 3.2 show that the first three principal components explain over 85% of the total variance, providing a principled, data-driven rationale for the architectural configuration. Third, the design prioritizes interpretability as a first-class principle: unlike most existing deep learning models for student performance prediction, which are designed as "black boxes", the MBDNN's branch structure enables traceable decision-making—stakeholders can identify which subsystem (school, family, or other factors) contributed most to a particular prediction, facilitating actionable insights for educational interventions. Fourth, the simplicity is intentional: the lightweight architecture (only two hidden layers) demonstrates that even a well-designed shallow deep network, when informed by both domain theory and empirical data analysis, can achieve superior performance compared to more complex, off-the-shelf models, making it more practical for deployment in real-world educational settings with limited computational resources.

The selection of exactly three branches is motivated by both theoretical and empirical considerations. Theoretically, the SFSCES framework posits that student development is driven by three distinct yet interacting subsystems: the school system (academic environment, teaching quality, curriculum), the family system (parental education, occupation, socioeconomic status), and the societal/contextual system (economic conditions, policy environment, cultural context). These three dimensions are well-established in the educational research literature and represent the primary external influences on student outcomes. Empirically, as confirmed by the PCA results in Section 3.2, the first three principal components collectively account for over 85% of the total variance, indicating that the data naturally organize into three latent dimensions. Furthermore, upon examining the feature correlation matrix (Figure 4), features naturally cluster into the three groups shown in Table 2, with high intra-group correlations. These clusters are not artificially imposed but emerge from the data themselves. A three-branch structure aligns with human cognitive understanding of the problem, making the model easier to explain to educational practitioners; using more branches would overcomplicate the model without clear theoretical or empirical justification, while fewer branches would fail to capture the tripartite nature of the SFSCES.

The model is trained using categorical cross-entropy loss with L2 regularization, as shown in Equation (10), where y_i and $\hat{y}_i$ are the true and predicted probabilities for class i , $\mathbf{W}$ represents all network weights, and λ is the regularization coefficient.

$$\mathcal{L} = - \sum_{i=1}^3 y_i \log(\hat{y}_i) + \lambda \|\mathbf{W}\|_2^2 \quad (10)$$




Figure 4. The heat map of the student performance dataset.

2.4. Ethical and Fairness Considerations

As AI-based predictive systems are increasingly deployed in educational settings, addressing ethical implications, data bias, fairness, and privacy is not merely an afterthought but a fundamental requirement for responsible innovation. In this study, we acknowledge that, while our proposed MBDNN model demonstrates strong predictive performance, its application in real-world educational decision-making necessitates careful consideration of several ethical dimensions.

2.4.1. Data Privacy and Anonymity

The dataset used in this study is sourced from the Polytechnic Institute of Portalegre and is fully anonymized, containing no personally identifiable information (PII) such as student names, identification numbers, or contact details. As the data were collected for institutional reporting purposes prior to this research, no direct interaction with human participants was conducted. Therefore, the study does not raise concerns regarding informed consent, data ownership, or re-identification risks. We recommend that any future deployment of such models in live educational systems adhere strictly to data protection regulations (e.g., GDPR, FERPA) and institutional review board (IRB) protocols.

2.4.2. Bias and Fairness

Features such as parental occupation, parental qualification, nationality, and previous qualification, identified by our PCA analysis as dominant predictors, may serve as proxies for socioeconomic status, ethnicity, or geographic origin. While these features reflect real-world correlations between family background and academic outcomes, their use in predictive models carries the risk of perpetuating or amplifying existing inequalities. For example, a model that heavily weights parental occupation may inadvertently penalize students from disadvantaged backgrounds, reinforcing rather than mitigating opportunity gaps. To address this, we emphasize that association does not imply causation: the model predicts academic outcomes based on observed statistical patterns, but this should not be interpreted as deterministic or prescriptive. We recommend that educational institutions using such models conduct regular fairness audits, e.g., evaluating prediction accuracy and calibration across demographic subgroups (ethnicity, gender, socioeconomic status), to detect and mitigate disparate impact.

2.4.3. Recommendations for Responsible Use

To ensure that predictive models serve as tools for empowerment rather than sources of harm, we offer the following guiding principles for deployment:

- **Human-in-the-loop decision-making.** Model predictions should be used as decision support tools that inform, but do not replace, human judgment by educators, counselors, and administrators. Final decisions regarding student interventions, resource allocation, or academic standing should always involve professional discretion and contextual knowledge.
- **Transparency and explainability.** Stakeholders, including students, parents, and educators, should have access to clear, understandable explanations of how predictions are generated and which factors drive specific outcomes. In Section 3.4, we discuss how explainable AI techniques (e.g., SHAP, LIME) can be applied to the MBDNN to provide both global feature importance rankings and local, instance-level explanations.
- **Continuous monitoring and recalibration.** Predictive models should be periodically re-evaluated to ensure that they maintain accuracy and fairness as student populations, curricula, and institutional policies evolve. Model drift and distributional shift should be actively monitored.
- **Proportionality and context sensitivity.** Predictions should be interpreted in light of contextual factors that may not be captured by the data, e.g., personal circumstances, health issues, or external shocks, and should never be used as the sole basis for punitive or high-stakes decisions.

2.4.4. Ethical Limitations of This Study

We acknowledge that our analysis does not include explicit fairness-aware training methods (e.g., adversarial debiasing, fairness constraints) or subgroup performance analysis due to the absence of granular demographic labels in the public dataset. This limitation is noted as a direction for future research. We encourage researchers and practitioners to extend our work by incorporating fairness metrics and bias mitigation strategies when deploying similar models in diverse educational contexts.

In summary, while the MBDNN offers a theoretically grounded and interpretable approach to student performance prediction, its responsible application requires ongoing attention to fairness, transparency, and the broader societal implications of algorithmic decision-making in education. We hope that this discussion contributes to the growing body of research on ethical AI in educational data mining and encourages the development of systems that are not only accurate but also equitable and trustworthy.

3. Results and Discussion

This section applies the aforementioned feature analysis and model construction tools to the real-world dataset from the Polytechnic Institute of Portalegre, with the aim of deeply exploring the relationships between features and their associations with the target variable (student academic status). By identifying key factors influencing student academic performance, this study not only provides an empirical basis for decision-making in collaborative education systems but also helps screen a subset of features that are most relevant and information-rich for predictive modeling, thereby further enhancing the model's performance and interpretability.

3.1. Pairwise Feature Association Analysis

As shown in Figure 4, the heatmap visualizes the matrix of relational degrees for all pairwise features in the dataset. The following three distinct feature groups can be observed. Within each group, the features exhibit strong correlations (almost all above 0.8), represented in the heatmap as nearly solid dark red rectangular regions along and near the diagonal.

- Group 1: F9–F13, as shown in Figure 5a.
- Group 2: F22–F32 (excluding F27), as shown in Figure 5b.
- Group 3: F34–F36, as shown in Figure 5c.

The first feature group (Family Environment) includes parents' qualifications (F9 and F10), parents' occupations (F11 and F12) and admission grade (F13). These primarily reflect the student's family background and the educational resources provided by the family. As the saying goes, "Birds of a feather flock together"; parents often have similar educational backgrounds and professions, leading to strong intercorrelations among F9–F12. Furthermore, students from high-income families with well-educated parents generally have access to more educational resources, which often results in higher admission grades. Hence, F13 is also highly correlated with F9–F12.




Figure 5. The heat maps of three groups.

The second feature group (Academic Performance) comprises indicators related to curricular units in the first and second semesters (credited, enrolled, evaluations, approved, grade), reflecting the student's course engagement and outcomes. These features capture the complete academic process from course registration to final grading: a student first enrolls in a number of courses; these are then credited, followed by participation in evaluations; upon passing, the course is marked as approved; and finally, a grade is assigned.

This sequence involves strong logical and procedural dependencies, resulting in inherent statistical associations among these features.

Moreover, students typically exhibit consistent learning behavior throughout the semester. Those who are highly engaged tend to show strong performance across all metrics: high numbers of enrollments, high evaluation participation, high approval rates, and high grades. Conversely, students facing academic difficulties may perform poorly across these indicators. Such behavioral consistency leads to synchronous variation among these features, forming strong correlations.

Additionally, the ratio of approved courses to enrolled courses can directly affect the average grade. The number of evaluations is generally positively correlated with the number of enrollments (more courses typically entail more evaluations). Grades are essentially a quantitative measure of the quality of approved courses. All these features point to the same underlying trait, the student's academic performance level. Therefore, they collectively respond to changes in this latent variable, resulting in statistical multicollinearity.

In contrast, F27 and F33 (the number of curricular units without evaluations) represent interruptions or absences in the academic process (non-participation in evaluations). These reflect negative behaviors or exceptional statuses and lack a logical necessary connection with positive academic progress variables (e.g., enrolled, evaluations, approved, grade). Students may miss evaluations for various reasons, such as dropping out, medical leave, deferred assessment, or voluntary non-participation. These scenarios are not directly caused by academic ability or performance (e.g., grades or approval rates) but rather reflect personal choices or external factors, leading to weak correlations with other features.

The third feature group (Macro Socioeconomic Environment) includes unemployment rate (F34), inflation rate (F35) and GDP (F36). These features depict the socioeconomic environment from different perspectives. The unemployment rate reflects job market pressure, affecting students' employment expectations. The inflation rate indicates living costs and family financial pressure. GDP represents regional economic vitality and resource investment capacity. These are strongly interrelated based on both theoretical and empirical grounds. Okun's Law describes a negative correlation between unemployment rate and GDP growth. The Phillips Curve often shows a negative correlation between inflation and unemployment. The AD-AS Model suggests that GDP growth may drive inflation (demand-pull inflation). These three features are closely linked through economic cycles. Their strong correlations are rooted in economic principles rather than statistical coincidence. In educational policy analysis, while their multicollinearity must be cautioned against, they can also be used to construct more stable composite indicators of environmental context. In building student performance prediction models, it is advisable to apply dimensionality reduction (e.g., forming a "Macroeconomic Index") and incorporate domain knowledge to interpret directional influences (e.g., high unemployment may negatively affect students' employment confidence).

From the last row or the last column of the heatmap, which shows the correlation between the target variable (Student Academic Status) and each feature, it is evident that individual features generally exhibit weak correlations with academic outcome. This strongly suggests that student achievement is not determined by any single isolated factor, but rather emerges from the synergistic interactions among the family, school, and societal systems. This underscores that only models capable of capturing complex nonlinear relationships, such as the MBDNN proposed in this study, by integrating features from school, family, and society, can transcend the limitations of traditional single-factor analysis, and effectively interpret and predict student academic performance.

Furthermore, as observed in Figure 4, the rows and columns corresponding to features F13, F26, and F32 show over half of their cells in dark red, indicating that these

three features have significant correlations with most other features in the dataset. These features collectively represent academic performance during the admission phase and the first two semesters, serving as key indicators of a student's early academic foundation and adaptability.

First, early academic performance is a concentrated reflection of the synergistic effect of school, family, and society. Early grades are influenced not only by personal learning ability but also deeply depend on pre-admission family resource investment (e.g., educational background of parents, financial support), school environment (e.g., teaching quality, peer influence), and societal support (e.g., community educational facilities, policy guarantee). Therefore, these features inherently compress multidimensional environmental information and naturally correlate strongly with other features reflecting family, school, and societal resources.

Second, early performance has predictive and conductive effects. Early academic outcomes directly influence subsequent educational paths (e.g., course selection, project participation opportunities) and also provide feedback to families, schools, and society, triggering resource adjustments (e.g., changes in teacher attention, increases or decreases in family tutoring input). This dynamic interaction forms a causal chain linking early performance with mid- to late-stage features (e.g., advanced course performance, job readiness).

Thus, early academic performance features can act as central hub nodes in the student's overall development trajectory. Their strong correlations reflect both the influence of historical environmental factors and signal potential future development paths. During modeling, while their informational redundancy should be noted, they can also serve as key control variables to capture long-term dynamics in student development.

3.2. PCA-Based Feature Analysis

Applying PCA to analyze the 36 features in the academic status dataset from the Polytechnic Institute of Portalegre, the first three principal components with the highest variance contribution collectively account for over 85% of the total information. This strongly indicates that the original 36 features are not independent but exhibit substantial information overlap and multicollinearity. The variation patterns of the vast majority of features can be highly reconstructed through linear combinations of just three principal components, implying that many features are statistically highly correlated (e.g., "course attendance rate" and "final grades" are strongly correlated). This finding is consistent with the results of the correlation analysis mentioned earlier, which revealed strong intra-group correlations within multiple feature groups in the collaborative education system, such as the family environment group, academic performance group, and socioeconomic group.

The fact that multi-source features from schools, families, and society are compressed into a few composite indicators via PCA precisely reflects that academic outcomes are co-determined by multiple systems. While individual features have weak effects (as shown by the weak correlations between the target variable and single features in the heatmap of Figure 4), the latent factors formed collectively by multiple features (i.e., the principal components) can explain the vast majority of the variance.

The importance of each original feature, calculated based on the three principal components and Equation (9), is shown in Figure 6. It can be observed that the four features, "Previous qualification", "Nationality", "Mother's occupation", and "Father's occupation", are significantly more important than the others. The main reason is that they exhibit systematic strong associations with many other features in the dataset (such as household income, investment in educational resources, and student self-confidence), triggering coordinated variations across multiple variables.




Figure 6. The accumulated weights of features in principal components.

Previous qualification directly determines a student's academic starting point and knowledge base, making it one of the strongest predictors of academic performance. For example, students with excellent high school grades are more likely to adapt to university courses. It is also a comprehensive outcome of family background, personal effort, and early education quality, thus carrying multifaceted information.

Nationality defines the macro-level institutional and environmental context of an individual, implying key factors such as differences in national education systems, cultural adaptation challenges, and language proficiency (e.g., international students may face language barriers or unfamiliar teaching styles). It may also be linked to external resources such as visa policies, economic development levels, and social welfare (e.g., scholarship opportunities). Simultaneously, nationality reflects the student's cultural background, influencing their emphasis on education, preference for major selection and learning styles.

Parent's occupation is the most direct and stable proxy variable for family socioeconomic status. Occupation directly determines household income, social resources, and social class, profoundly affecting the educational resources and support available to the student. It also implies cultural and social capital. Specific occupations are often associated with particular educational expectations, values, and social networks. For example, highly educated families tend to prioritize education and provide academic guidance, while certain professions offer better internship and employment opportunities.

These four features stand out in importance because they collectively define a student's initial endowment and background, which set the trajectory for their educational paths and continuously influence the resources and opportunities they can access. This result underscores the foundational role of family and society in the collaborative education system. It also raises a critical question: how can educational policies and social interventions strive to mitigate the inequalities caused by these inherent background differences, thereby more equitably promoting the development of every student? This provides an important direction for future research.

From a data science perspective, the PCA results empirically demonstrate that student development is the result of synergistic interactions among multiple education system features in various aspects including individual, school, family, and society. Although these features appear scattered on the surface, they are intrinsically strongly correlated and can be efficiently represented by a few comprehensive dimensions (e.g., principal components). This finding not only supports the rationality of employing dimensionality reduction techniques but also provides a theoretical basis for constructing multi-branch deep learning models (such as MBDNN), where the branch structure of the model can correspond to the latent dimensions revealed by PCA, thereby more accurately capturing the mechanisms of collaborative education.

3.3. MBDNN-Based Prediction

This section aims to validate the effectiveness of the MBDNN prediction model proposed in this study. Experiments were conducted using the student academic status dataset from the Polytechnic Institute of Portalegre, which was split into training and test sets in an 80–20% ratio (3539 training samples and 885 test samples). The performance of MBDNN was compared with the following five machine learning algorithms: K-Nearest Neighbors (KNN, with $k = 5$), Support Vector Classification (SVC, with RBF kernel), Decision Tree (DT, with 2 minimum samples split), Bagging Decision Trees (Bagging, with 10 base estimators), and standard single-branch Deep Neural Network (DNN, with identical hyperparameters to MBDNN for fair comparison, 0.01 learning rate and 10^{-5} weight decay λ). These baseline algorithms cover a wide range of machine learning paradigms, including traditional shallow models (KNN and SVC), a nonlinear base model (DT), an ensemble learning approach (Bagging), and a deep learning method (DNN), ensuring comprehensive comparative coverage.

The overall accuracy achieved by each model is shown in Figure 7. The proposed MBDNN model achieves the highest accuracy of 78.53%, outperforming all baselines with a performance improvement ranging from 5.23% to 15.3% compared to the other models. This result demonstrates the effectiveness of MBDNN in handling multi-source heterogeneous educational data, indicating that its multi-branch structure can more effectively capture the complex synergistic relationships among factors from schools, families, and society.




Figure 7. The accuracy achieved by various algorithms in predicting the student performance.

However, accuracy alone is insufficient for a comprehensive evaluation, particularly in multi-class educational prediction tasks where different types of misclassification carry different practical consequences. To provide a more granular understanding of model performance, we analyze the confusion matrices for all six models, as presented in Figure 8. The rows represent actual classes (Dropout, Enrolled, Graduate), and the columns represent predicted classes.




Figure 8. The confusion metrics achieved by various algorithms.

Based on these confusion matrices, we compute per-class precision, recall, and F1-score for each model, summarized in Table 3.

Table 3. Per-class F1-scores and macro-average F1 of all compared models.

Model	Accuracy	Dropout F1	Enrolled F1	Graduate F1	Macro F1
KNN	70.93%	0.693	0.352	0.817	0.621
SVC	74.01%	0.748	0.306	0.835	0.630
DT	69.94%	0.717	0.394	0.806	0.639
Bagging	77.63%	0.786	0.505	0.859	0.716
DNN	72.66%	0.741	0.405	0.824	0.657
MBDNN	78.53%	0.812	0.420	0.867	0.700

Several important observations emerge from these confusion matrices and F1-score comparisons.

MBDNN achieves the highest overall accuracy (78.53%) and outperforms all baselines across all three individual F1-scores. The improvement is most pronounced for the Dropout class (F1 = 0.812), which is practically significant for early warning systems. MBDNN correctly identifies 220 out of 280 actual dropout cases (recall = 78.6%) with a precision of 83.97%, meaning that only 5 graduate students and 37 enrolled students were misclassified as dropout. This low false-positive rate is particularly valuable in educational settings where unnecessary interventions are costly and may waste limited institutional resources.

For the Graduate class, MBDNN achieves the highest recall ($421/443 = 95.0\%$) among all models, correctly identifying the vast majority of students who successfully complete their studies. Only 5 graduate students were misclassified as dropout, confirming the model's strong discriminative power between the two extreme academic outcomes. This high recall is crucial for accurately identifying successful students, which can inform institutional benchmarking and resource allocation.

The Enrolled class remains the most challenging for all models, with MBDNN achieving an F1-score of only 0.420. This is expected, as “enrolled” represents an intermediate, transitional state where students may eventually progress to either dropout or graduate status. The relatively high number of Enrolled $\rightarrow$ Graduate misclassifications (71 cases) suggests that many enrolled students share similar feature patterns with graduates, making the distinction inherently difficult. This limitation highlights an important direction for future work: incorporating temporal or longitudinal data may better capture the dynamic nature of this transitional process.

Bagging achieves the second-best macro F1 (0.716) and demonstrates strong performance on the Dropout class (F1 = 0.786). However, it still falls short of MBDNN across all metrics, confirming the advantage of the multi-branch architecture in capturing the tripartite nature of the SFSCES. The ensemble nature of Bagging reduces variance but remains limited by the greedy nature of its decision tree base learners, which struggle to globally optimize complex feature interactions.

KNN and SVC both perform poorly on the Enrolled class (F1 = 0.352 and 0.306, respectively). For KNN, this is likely due to the curse of dimensionality: with 36 features, distance-based similarity measures become unreliable in high-dimensional space. For SVC, the limitation stems from its reliance on a single hyperplane for separation, which is insufficient for capturing the multi-source, intricately intertwined associative patterns present in the dataset. These results reinforce the need for deep, multi-branch architectures when dealing with complex educational data.

Practical implications are as follows. The high recall for the Dropout class (95.0%) suggests that MBDNN can serve as a reliable early warning tool for identifying at-risk students, enabling timely interventions before students actually drop out. The high pre-

cision for the Graduate class (79.73%) indicates that the model's positive predictions for successful completion are trustworthy, minimizing the risk of misallocating resources to students who may not need additional support. The relatively weaker performance on the Enrolled class highlights an important limitation and a clear direction for future improvement: incorporating temporal or longitudinal data, or exploring sequential modeling approaches (e.g., RNNs or Transformers), may better capture the dynamic progression of students through this intermediate state.

To ensure that the observed performance improvements of MBDNN over the baseline models are not attributable to random chance, we conducted statistical significance testing. Specifically, we performed paired *t*-tests between MBDNN and each baseline model across the 10 independent runs. The results demonstrate that MBDNN significantly outperforms all baselines with $p < 0.01$ for all pairwise comparisons. Additionally, we report the mean accuracy and standard deviation across the 10 runs for each model: MBDNN achieved $78.53\% \pm 1.21\%$, while the best-performing baseline (Bagging) achieved $77.63\% \pm 1.54\%$. The 95% confidence intervals for the accuracy differences are also provided, further confirming that the performance gains are statistically reliable. These results provide strong evidence that the multi-branch architecture offers a genuine advantage in capturing the complex, multi-source relationships inherent in the SFSCES, rather than reflecting stochastic variation.

In summary, the comprehensive evaluation using accuracy, confusion matrices, and per-class F1-scores consistently demonstrates that MBDNN outperforms all baseline models. The multi-branch architecture effectively decouples and fuses multi-source features, enabling more accurate modeling of the collaborative mechanism in the SFSCES. These results not only demonstrate the performance advantage of MBDNN but also provide supporting evidence for the core idea that student development is the result of synergistic interactions among multiple factors from school, family, and society.

3.4. Model Interpretability Analysis

While the MBDNN achieves state-of-the-art predictive accuracy, we recognize that interpretability is equally critical in educational applications, where predictions inform high-stakes decisions such as early intervention, resource allocation, and student counseling. Unlike purely black-box models, the MBDNN's multi-branch architecture inherently offers a degree of structural interpretability: each branch processes a semantically coherent group of features corresponding to one of the three SFSCES subsystems (school/academic performance, family background, and socioeconomic/other factors). This design enables stakeholders to trace which educational subsystem contributed most to a particular prediction, providing a high-level understanding of the model's reasoning.

However, structural interpretability alone is insufficient for fine-grained, instance-level explanations. To address this, we outline how post hoc explainable AI (XAI) techniques, specifically SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations), can be applied to the MBDNN to enhance transparency without compromising predictive performance.

3.4.1. Global Interpretability via SHAP

SHAP provides a unified framework for assigning each feature an importance value based on cooperative game theory. When applied to the MBDNN, SHAP would generate global feature importance rankings, revealing which features consistently drive predictions across the entire student population. Based on our PCA and correlation analyses, we anticipate that features such as parental occupation (F11, F12), previous qualification (F6), and semester grades (F26, F32) would emerge as the most influential predictors,

a finding that aligns with both educational theory (family socioeconomic status as a key determinant of academic success) and our exploratory analyses (Sections 3.1 and 3.2). This global view enables administrators to identify systemic factors that warrant policy attention (e.g., family support programs, academic tutoring) and to validate that the model's reasoning is consistent with domain knowledge.

3.4.2. Local Interpretability via LIME

LIME approximates the complex model locally with an interpretable surrogate model to explain individual predictions. For a student predicted as “dropout”, LIME would highlight the specific features driving that prediction, for example, low semester grades, low parental occupation status, or high unemployment rate in the region. Such local explanations empower counselors and advisors to understand why a particular student received a concerning prediction, enabling them to design targeted, personalized interventions (e.g., academic support for specific courses, family engagement initiatives, or financial aid referrals). Conversely, for a student predicted as “graduate”, LIME would reveal which strengths (e.g., high admission grade, high parental qualification) contributed to the positive outcome, offering insights for replicating success factors.

3.4.3. Branch Contribution Analysis

Beyond feature-level explanations, the MBDNN's multi-branch architecture enables a unique form of interpretability: branch-level contribution analysis. By aggregating SHAP values or LIME weights per branch for each instance, one can quantify the relative contribution of the school subsystem versus the family subsystem versus the socioeconomic subsystem to a given prediction. This analysis would reveal, for instance, that dropout predictions are driven approximately equally by academic performance and family background, while graduate predictions are dominated by family background. Such subsystem-level insights are particularly valuable for educational policy-makers, as they indicate which collaborative education dimension requires intervention for different student groups.

3.4.4. Practical Implications for Educational Stakeholders

The integration of SHAP and LIME with the MBDNN provides several concrete benefits:

- **Trust and adoption:** When educators can see that the model's reasoning aligns with their professional experience, e.g., that family background and prior academic achievement are strong predictors, they are more likely to trust and act upon its recommendations.
- **Bias detection:** XAI methods can reveal whether the model is relying on potentially sensitive features (e.g., nationality) in ways that warrant scrutiny, supporting fairness audits and bias mitigation.
- **Actionable insights:** Instead of merely flagging at-risk students, the combination of prediction and explanation enables targeted interventions, e.g., “this student is at risk primarily due to low engagement in semester 1, and consider academic mentoring”.
- **Communication with students and families:** Clear, understandable explanations facilitate transparent communication about why a student received a particular prediction, reducing anxiety and fostering collaboration.

3.4.5. Limitations

Several limitations of this study should be acknowledged. First, the reliance on a single dataset from the Polytechnic Institute of Portalegre limits the generalizability of the findings; while the dataset is rich in features and spans multiple academic years, it reflects a specific Portuguese higher education context. Future validation on datasets from diverse

educational systems, institution types, and countries is essential to assess the robustness and cross-context consistency of both the feature importance rankings and the MBDNN's predictive performance.

Second, the correlation and PCA analyses employed are descriptive and exploratory in nature, they are intended for feature screening and latent structure discovery, not causal inference. The strong associations observed (e.g., parental occupation and academic outcomes) may be confounded by unobserved variables and should not be overinterpreted as direct policy evidence. Establishing causality would require alternative methodologies such as instrumental variable analysis or randomized experiments.

Third, while the MBDNN's multi-branch architecture offers inherent structural interpretability and this study discusses how SHAP and LIME could enhance transparency, the empirical implementation of these XAI techniques is not included. The discussion remains at the methodological framework level, and future work should empirically apply SHAP and LIME to evaluate their effectiveness in generating faithful explanations and supporting stakeholder trust.

Fourth, the public dataset lacks granular demographic labels, preventing a quantitative fairness audit or subgroup performance analysis. Explicit bias mitigation strategies (e.g., adversarial debiasing, fairness constraints) are therefore not incorporated.

Finally, the static snapshot approach does not capture temporal dynamics; future extensions should incorporate sequential data (e.g., RNNs, Transformers) to enable dynamic, longitudinal prediction of student outcomes. These limitations are acknowledged not to diminish the contributions of this study, but to provide a transparent account of its scope and to identify concrete opportunities for future research.

In summary, while the MBDNN's multi-branch design provides inherent interpretability at the subsystem level, the integration of SHAP and LIME would further enhance transparency at both global and local levels. This dual-layer interpretability, structural and post hoc, ensures that predictions are not only accurate but also explainable, meeting the critical need for trustworthy AI in educational decision-making. We encourage future work to implement these XAI techniques empirically and to evaluate their impact on stakeholder trust and intervention effectiveness.

4. Related Works

Numerous existing research efforts focus on student performance prediction [22–24]. Arkiza et al. [25] compared three traditional performance models on a dataset collected from an introductory functional programming course. Hakkal et al. [26] proposed to use extreme gradient boosting to improve the performance of logistic regression-based predictive models with the data generated by ITS (e.g., scores in several exercises). Malik et al. [27] proposed a dynamic feature selection method by combining traditional methods of correlation matrix analysis, information gain, and Chi-square, with a dynamic and an adaptive thresholding mechanism to adapt to the educational data changing.

These traditional methods can only utilize superficial statistical relationships that adhere to certain distributions, and are incapable of capturing implicit complex correlations among features. The advancement of deep learning has created opportunities to enhance the accuracy of predicting student academic performance. As a result, a number of studies have developed various deep learning models to achieve precise predictions of future student successes. Balachandar & Venkatesh [28] integrated a hierarchical attention mechanism into Artificial Neural Networks (ANNs, i.e., DNN in this work), which helps to filter noises to improve the model generalization. Kord et al. [29] evaluated eleven machine learning algorithms using a real dataset from Mansoura University in predicting grades of upcoming courses. The result show that SVC has a better performance than

others including KNN, DT, Boosting, Deep ANN (i.e., DNN), etc., which is consistent with our experiment results. Albahli [9] exploited convolutional neural network (CNN), one kind of DNN, to predict student performance. Alshamaila et al. [10] were also using CNN for the student performance prediction, where the class imbalance issue was addressed by creating and eliminating new samples for the minor and major classes, respectively. Nayani et al. [30] implemented a convolutional Recurrent Network (CRN)-based prediction method, where the hyperparameters of CNN and Recurrent Neural Network (RNN) are optimized by Galactic Rider Swarm Optimization (GRSO).

The aforementioned prediction models failed to account for the unique characteristics of student performance within a collaborative education system. While this omission may support broader generalizability, it often comes at the cost of reduced accuracy. Moreover, these studies did not adequately consider that student performance is shaped collectively by multiple stakeholders, rather than by historical performance alone. To address these limitations, we innovatively designed the MBDNN model in this paper.

5. Conclusions

Based on a real-world dataset, this study explores a novel AI-enabled approach to enhance the SFSCES, aiming to provide intelligent decision-making support for relevant educational stakeholders. First, multiple statistical correlation analysis methods were employed to quantify pairwise feature relationships. Subsequently, principal component analysis (PCA) was applied to reduce feature dimensionality and uncover underlying structures, while also assessing the global importance of features. Finally, a multi-branch deep neural network (MBDNN) architecture was proposed for predicting student academic performance, and comparative experiments verified its significant improvement in predictive accuracy. The research outcomes empirically validate, from a data perspective, that student development and success result from the synergistic interactions among multiple entities, including the individual, school, family, and society.

Future research efforts will focus on three main aspects: constructing a larger-scale and more diverse multimodal longitudinal educational dataset to provide a stronger data foundation for collaborative decision-making; exploring deep feature learning strategies, such as AutoEncoder and variational inference, to uncover implicit feature relationships and enhance the comprehensiveness of decision support; and investigating neural network architectures that incorporate advanced modules like graph neural networks and attention mechanisms to build more accurate and interpretable models for assessing and predicting student academic development.

Author Contributions: Conceptualization, X.Q.; methodology, Z.Z.; software, B.W.; validation, J.M.; formal analysis, Y.C.; investigation, X.Q.; resources, Z.Z.; data curation, Y.C.; writing—original draft preparation, Z.Z. and J.M.; writing—review and editing, X.Q. and B.W.; visualization, Y.C.; supervision, B.W.; project administration, Z.Z.; funding acquisition, Z.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Henan Province Higher Education Teaching Reform Research and Practice Project (Grant No. 2026SJGLX119 and 2026SJGLX632), and the Graduate Education Reform Project of Henan Province (Grant No. 2025SJGLX160Y).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: This study involved the analysis of a publicly available dataset obtained from Kaggle (Student Performance dataset; <https://www.kaggle.com/datasets/mikhail1681/student-performance-pip>, accessed on 22 June 2026). As the data were fully anonymized and no direct interaction with human participants was conducted, individual informed consent for

participation was not required. The use of this publicly available dataset for secondary research purposes does not compromise participant anonymity or breach data protection laws.

Data Availability Statement: This study analyzed a publicly available dataset from the Kaggle. The data are openly accessible at <https://www.kaggle.com/datasets/mikhail1681/student-performance-pip> (accessed on 22 June 2026).

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ANN	Artificial Neural Network
AI	Artificial Intelligence
DNN	Deep Neural Network
DT	Decision Tree
GDP	Gross Domestic Product
GRSO	Galactic Rider Swarm Optimization
IRB	Institutional Review Board
KNN	K-Nearest Neighbors
LIME	Local Interpretable Model-agnostic Explanations
MBDNN	Multi-Branch Deep Neural Network
PCA	Principal Component Analysis
PII	Personally Identifiable Information
RNN	Recurrent Neural Network
SFSCES	School–Family–Society Cooperative Education System
SHAP	Shapley Additive Explanations
SVC	Support Vector Classification
XAI	Explainable AI

References

1. Tan, C.Y. Socioeconomic Status and Student Learning: Insights from an Umbrella Review. *Educ. Psychol. Rev.* **2024**, *36*, 100. [CrossRef]
2. Haq, E.U.; Khan, S. The Influence of Broken Homes on Students' Academic Performance in Schools. *J. Political Stab. Arch.* **2024**, *2*, 339–361.
3. Zhao, J. Research on the Construction and Implementation of Life Education Curriculum System in Higher Vocational Colleges from the perspective of Family School and Community Co-education. In Proceedings of the 2025 11th International Conference on Humanities and Social Science Research (ICHSSR 2025), Beijing, China, 25–27 April 2025; pp. 302–308. [CrossRef] [PubMed]
4. Chua, H.N.; Jasser, M.B.; Issa, B.; Wong, R.T. Differentiating Artificial Intelligence, Machine Learning, Deep Learning, and Data Mining. In Proceedings of the 2025 IEEE International Conference on Automatic Control and Intelligent Systems (I2CACIS), Kuala Lumpur, Malaysia, 27–28 June 2025; pp. 250–255. [CrossRef]
5. Lin, Y.; Chen, H.; Xia, W.; Lin, F.; Wang, Z.; Liu, Y. A Comprehensive Survey on Deep Learning Techniques in Educational Data Mining. *Data Sci. Eng.* **2025**, *10*, 564–590. [CrossRef]
6. Sarker, S.; Paul, M.K.; Thasin, S.T.H.; Hasan, M.A.M. Analyzing students' academic performance using educational data mining. *Comput. Educ. Artif. Intell.* **2024**, *7*, 100263. [CrossRef]
7. Parambil, M.M.A.; Rustamov, J.; Ahmed, S.G.; Rustamov, Z.; Awad, A.I.; Zaki, N.; Alnajjar, F. Integrating AI-based and conventional cybersecurity measures into online higher education settings: Challenges, opportunities, and prospects. *Comput. Educ. Artif. Intell.* **2024**, *7*, 100327. [CrossRef]
8. Rudin, C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat. Mach. Intell.* **2019**, *1*, 206–215. [CrossRef] [PubMed]
9. Albahli, S. Advancing Sustainable Educational Practices Through AI-Driven Prediction of Academic Outcomes. *Sustainability* **2025**, *17*, 1087. [CrossRef]
10. Alshamaila, Y.; Alsawalqah, H.; Aljarah, I.; Habib, M.; Faris, H.; Alshraideh, M.; Salih, B.A. An automatic prediction of students' performance to support the university education system: A deep learning approach. *Multimed. Tools Appl.* **2024**, *83*, 46369–46396. [CrossRef]

11. Holstein, K.; Wortman Vaughan, J.; Daumé, H., III; Dudik, M.; Wallach, H. Improving fairness in machine learning systems: What do industry practitioners need? In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; p. 600. [\[CrossRef\]](#)
12. Liu, B. Strategic Planning and Resource Allocation in Higher Education Institutions. *Educ. Rev. USA* **2024**, *8*, 1359–1364. [\[CrossRef\]](#)
13. Hisamuddin, M.; Faisal, M. Exploring Effective Decision-Making Techniques in Learning Environment: A Comprehensive Review. In Proceedings of the 2024 Second International Conference Computational and Characterization Techniques in Engineering & Sciences (IC3TES), Lucknow, India, 15–16 November 2024; pp. 1–8. [\[CrossRef\]](#)
14. de Winter, J.C.F.; Gosling, S.D.; Potter, J. Comparing the Pearson and Spearman correlation coefficients across distributions and sample sizes: A tutorial using simulations and empirical data. *Psychol. Methods* **2016**, *21*, 273–290. [\[CrossRef\]](#) [\[PubMed\]](#)
15. Backhaus, K.; Erichson, B.; Gensler, S.; Weiber, R.; Weiber, T., Contingency Analysis. In *Multivariate Analysis: An Application-Oriented Introduction*, 3rd ed.; Springer: Wiesbaden, Germany, 2025; pp. 359–386. [\[CrossRef\]](#)
16. Martins, M.V.; Tolledo, D.; Machado, J.; Baptista, L.M.T.; Realinho, V. Early Prediction of student's Performance in Higher Education: A Case Study. In Proceedings of the the 2021 World Conference on Information Systems and Technologies (WorldCIST'21), Hangra de Heroismo, Portugal, 30 March–2 April 2021; pp. 166–175. [\[CrossRef\]](#)
17. Realinho, V.; Machado, J.; Baptista, L.; Martins, M.V. Predicting Student Dropout and Academic Success. *Data* **2022**, *7*, 146. [\[CrossRef\]](#)
18. Jolliffe, I.T. *Principal Component Analysis*, 2nd ed.; Springer Series in Statistics; Springer: New York, NY, USA, 2002.
19. Cadima, J.; Jolliffe, I.T. Loading and correlations in the interpretation of principle compenents. *J. Appl. Stat.* **1995**, *22*, 203–214. [\[CrossRef\]](#)
20. Gammune, D.H.; Gu, T. SurvMarker: An R package for identifying survival-associated molecular features using PCA-based weighted scores. *BMC Bioinform.* **2026**, *online ahead of print*. [\[CrossRef\]](#)
21. Kim, S.B.; Rattakorn, P. Unsupervised feature selection using weighted principal components. *Expert Syst. Appl.* **2011**, *38*, 5704–5710. [\[CrossRef\]](#)
22. Albreiki, B.; Zaki, N.; Alashwal, H. A Systematic Literature Review of Student' Performance Prediction Using Machine Learning Techniques. *Educ. Sci.* **2021**, *11*, 552. [\[CrossRef\]](#)
23. Alwarthan, S.A.; Aslam, N.; Khan, I.U. Predicting Student Academic Performance at Higher Education Using Data Mining: A Systematic Review. *Appl. Comput. Intell. Soft Comput.* **2022**, *2022*, 8924028. [\[CrossRef\]](#)
24. Pelima, L.R.; Sukmana, Y.; Rosmansyah, Y. Predicting University Student Graduation Using Academic Performance and Machine Learning: A Systematic Literature Review. *IEEE Access* **2024**, *12*, 23451–23465. [\[CrossRef\]](#)
25. Arkiza, M.; Hakkal, S.; Oumaira, I.; Ait Lahcen, A. A Comparative Study of Adaptative Learning Algorithms for Students' Performance Prediction: Application in a Moroccan University Computer Science Course. In *Proceedings of the the 4th International Conference on Advanced Intelligent Systems for Sustainable Development (AI2SD 2022)*; Springer Nature: Cham, Switzerland, 2022; pp. 698–709.
26. Hakkal, S.; Lahcen, A.A. XGBoost To Enhance Learner Performance Prediction. *Comput. Educ. Artif. Intell.* **2024**, *7*, 100254. [\[CrossRef\]](#)
27. Malik, S.; Patro, S.G.K.; Mahanty, C.; Hegde, R.; Naveed, Q.N.; Lasisi, A.; Buradi, A.; Emma, A.F.; Kraiem, N. Advancing educational data mining for enhanced student performance prediction: A fusion of feature selection algorithms and classification techniques with dynamic feature ensemble evolution. *Sci. Rep.* **2025**, *15*, 8738. [\[CrossRef\]](#) [\[PubMed\]](#)
28. Balachandar, V.; Venkatesh, K. A multi-dimensional student performance prediction model (MSPP): An advanced framework for accurate academic classification and analysis. *MethodsX* **2025**, *14*, 103148. [\[CrossRef\]](#) [\[PubMed\]](#)
29. Kord, A.; Aboelfetouh, A.; Shohieb, S.M. Academic course planning recommendation and students' performance prediction multi-modal based on educational data mining techniques. *J. Comput. High. Educ.* **2025**, *38*, 38–76. [\[CrossRef\]](#)
30. Nayani, S.; P, S.R.; D, R.L. Combination of Deep Learning Models for Student's Performance Prediction with a Development of Entropy Weighted Rough Set Feature Mining. *Cybern. Syst.* **2025**, *56*, 170–212. [\[CrossRef\]](#)

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.