

The Relationship between the Use of Artificial Intelligence and the Quality of Learning in Higher Education Institutions: A Study Based on Predictive Models

¹Roberto Johan Barragan-Monrroy, ²Mario Javier Bonilla-Loor, ³Felipe Aguirre Chávez, ⁴Galo Enrique Vásquez Vásquez

¹Quevedo State Technical University, Quevedo, Ecuador

Email: roberto.barragan2014@uteq.edu.ec

ORCID: <https://orcid.org/1376-1388-1376-1388>

²Universidad Técnica de Manabí, Facultad de Agrocencias, Departamento de Procesos Agroindustriales, Portoviejo, Manabí, Ecuador

Email: mario.bonilla@utm.edu.ec

ORCID: <https://orcid.org/1376-1388-1376-1388>

³Universidad Marcelino Champagnat, Lima, Perú

Email: faguirre@umch.edu.pe

ORCID: <https://orcid.org/0000-0003-3474-2603>

⁴Universidad de las Fuerzas Armadas ESPE, Ecuador

Email: gevasquez5@espe.edu.ec

ORCID: <https://orcid.org/1376-1388-1376-1388>

HOW TO CITE:

Roberto Johan Barragan-Monrroy, Mario Javier Bonilla-Loor, Felipe Aguirre Chávez, Galo Enrique Vásquez Vásquez, (2026). The Relationship between the Use of Artificial Intelligence and the Quality of Learning in Higher Education Institutions: A Study Based on Predictive Models. International Journal of Special Education, 41(17s), 1376-1388

ABSTRACT:

The accelerated adoption of artificial intelligence (AI) in higher education has intensified the debate on whether AI tools genuinely enhance the quality of student learning or merely displace effort. This study examines the relationship between the intensity of AI use and the quality of learning in higher education institutions (HEIs) and, using a predictive-modeling approach, identifies which factors most strongly forecast learning quality. A quantitative, non-experimental, and predictive design was applied to a simulated cohort of 1,200 undergraduate students whose data-generating process was calibrated to effect sizes and relationships reported in the recent peer-reviewed literature. Nine predictors—AI use intensity, AI literacy, self-regulated learning (SRL), pedagogical guidance, engagement, prior GPA, study hours, socioeconomic index, and age—were modeled against a composite quality-of-learning outcome. Five algorithms were benchmarked: multiple linear regression, random forest, gradient boosting, extreme gradient boosting (XGBoost), and an artificial neural network (ANN). Multiple linear regression achieved the strongest generalization ($R^2 = 0.567$; RMSE = 6.04; 10-fold cross-validated $R^2 = 0.579$), followed closely by ensemble methods. Permutation importance identified SRL (32.2%), engagement (21.3%), and prior GPA (20.1%) as the dominant predictors, whereas raw AI use intensity contributed modestly (2.5%). Crucially, a significant AI-use $\times$ pedagogical-guidance interaction ($\beta = 0.45$, $p < 0.001$) indicated that AI use predicts higher learning quality only when embedded in structured pedagogical scaffolding. The findings reframe AI as an amplifier of sound learning conditions rather than an autonomous driver of quality, and they provide HEIs with an interpretable early-warning framework for data-informed decision-making.

COPYRIGHT STATEMENT:

Copyright: © 2026 Authors.
Open access publication under the terms and conditions of the Creative Commons Attribution (CC BY)

license (<http://creativecommons.org/licenses/by/4.0/>).

Keywords: artificial intelligence; higher education; quality of learning; predictive models; machine learning; learning analytics; self-regulated learning; academic performance

1. Introduction

Artificial intelligence (AI) has moved from the periphery to the center of higher education within a remarkably short period. Systematic evidence shows that AI-in-education publications rose two- to threefold in 2021–2022 relative to previous years, with applications clustering around five recurrent functions: assessment and evaluation, prediction, AI assistants, intelligent tutoring systems, and the management of student learning (Crompton & Burke, 2023). The public release of generative conversational agents accelerated this trajectory further, prompting institutions to reconsider assessment design, academic integrity policies, and the very definition of learning quality (Chaudhry et al., 2023). What was once a specialized research niche is now a mainstream concern for educators, administrators, and policymakers alike.

Yet the central empirical question remains unsettled: does the use of AI improve the quality of student learning, and if so, under what conditions? The literature is strikingly heterogeneous. A meta-analysis of 65 independent studies found that generative AI produced a medium positive effect on college students' academic achievement (Hedges' $g = 0.533$; Sun & Zhou, 2024). In contrast, quasi-experimental and mixed-methods work indicates that unguided AI use does not consistently improve—and may even depress—performance on higher-order tasks, while structured, pedagogically framed use yields substantial gains (Hassan, 2026). Some survey-based studies report only weak or non-significant associations between AI usage and learning outcomes once institutional readiness and student motivation are accounted for (Aslam et al., 2025). This divergence suggests that the effect of AI on learning is conditional rather than universal, and that aggregate main-effect estimates may obscure important moderators.

A parallel body of research has demonstrated the predictive value of AI itself for modeling academic outcomes. Machine-learning models trained on demographic, behavioral, and academic features can forecast student performance with useful accuracy, enabling the early identification of at-risk learners and targeted intervention (Yağcı, 2022; Ahmed et al., 2025; Guevara-Reyes et al., 2025). Increasingly, scholars argue that the frontier lies not in maximizing predictive accuracy in isolation, but in integrating

prediction with learning analytics to close the loop between forecasting and pedagogical action (Ouyang et al., 2023). Predictive modeling thereby serves a dual role in the present study: it is both the analytical lens through which the AI–learning relationship is examined and a practical instrument HEIs can deploy for decision support.

Despite this momentum, three gaps persist. First, much of the empirical work treats “AI use” as a monolithic main effect, seldom testing the moderators—such as pedagogical guidance and self-regulation—that theory predicts should govern its impact (Shi et al., 2025). Second, studies that report very high predictive accuracy (e.g., R^2 or accuracy above 0.95) frequently lack interpretability, limiting their usefulness for educational decision-making (Guevara-Reyes et al., 2025). Third, comparatively few studies foreground the construct of learning quality—as distinct from raw grade attainment—and situate AI use within a broader nomological network that includes engagement, prior achievement, and learner characteristics.

This study addresses these gaps by pursuing two objectives. The first is explanatory: to estimate the relationship between AI use intensity and the quality of learning while explicitly testing whether pedagogical guidance and self-regulated learning moderate that relationship. The second is predictive: to benchmark a set of interpretable and ensemble machine-learning models on their ability to forecast learning quality, and to identify—through permutation-based feature importance—which factors most strongly drive it. In doing so, the paper contributes an interpretable, theory-driven predictive framework that HEIs can adopt for early warning and resource allocation, and it advances a more nuanced account of AI as an amplifier of favorable learning conditions rather than an autonomous cause of quality.

The remainder of the article is organized as follows. Section 2 develops the theoretical background and hypotheses. Section 3 details the research design, data, measures, models, and evaluation protocol. Section 4 reports the descriptive, inferential, and predictive results. Section 5 discusses the findings in light of prior work and derives theoretical and practical implications. Section 6 acknowledges limitations and outlines future research, and Section 7 concludes.

2. Theoretical Background and Hypotheses

2.1 Artificial Intelligence in Higher Education

AI in education (AIEd) encompasses a family of technologies—machine learning, natural language processing, intelligent tutoring, and learning analytics—that emulate aspects of human cognition to support instruction, learning, and administration (Chen et al., 2020). Early systematic mapping warned that educators were underrepresented in AIEd research, which was dominated by technical disciplines and risked overlooking pedagogical grounding (Zawacki-Richter et al., 2019). More recent reviews document both a disciplinary rebalancing toward education departments and a geographical shift in research output, alongside a proliferation of tools aimed primarily at students (Crompton & Burke, 2023).

The mechanisms through which AI is theorized to benefit learning are well rehearsed: personalization of content and pace, immediate and targeted feedback, adaptive scaffolding, and the automation of routine assessment, which frees instructor time for higher-value interaction (Chen et al., 2020; Slimi, 2023). Systematic reviews of AI in personal and collaborative learning environments report gains in personalized learning, communication, engagement, and motivation, while consistently flagging ethical risks around privacy, bias, and over-reliance (Msambwa et al., 2025). Acceptance-oriented studies, grounded in the Unified Theory of Acceptance and Use of Technology, find that performance expectancy, effort expectancy, information accuracy, and pedagogical fit predict adoption and, through it, satisfaction and learning outcomes (Dahri et al., 2024; Venkatesh et al., 2003).

However, the evidence is not uniformly positive. AI-integrated applications can impose rigid frameworks that constrain creative thinking, heighten performance anxiety, and produce emotional disengagement when interactions feel impersonal (Lin et al., 2024). Empirical comparisons across pre-AI, unguided-AI, and structured-AI periods reveal that the sign and magnitude of AI's effect depend heavily on how its use is framed pedagogically (Hassan, 2026). These tensions motivate a conditional, moderator-sensitive account of AI's relationship with learning quality.

2.2 Conceptualizing the Quality of Learning

Quality of learning is a multidimensional construct that extends beyond grade attainment to encompass depth of understanding, engagement, critical thinking, self-regulation, and the transfer of knowledge to novel

tasks. In this study, learning quality is operationalized as a composite reflecting mastery, sustained engagement, and cognitive investment rather than a single summative grade. This framing aligns with evidence that AI's benefits are most visible in engagement, self-regulated learning, and higher-order skills, and least reliable when reduced to raw performance on summative assessments (Hassan, 2026; Anghel et al., 2025).

Two learner-level constructs are theoretically central. Self-regulated learning (SRL)—the capacity to plan, monitor, and adjust one's own learning—has long been associated with academic success (Zimmerman, 2002) and, in AI-supported contexts, has been shown to predict writing performance more strongly than AI literacy itself (Shi et al., 2025). Engagement, in turn, functions as a proximal mediator: structured AI use increases independent study engagement, which then translates into improved critical thinking, communication, and problem-solving (Hassan, 2026). These constructs anchor the expectation that AI use influences learning quality partly by amplifying self-regulation and engagement rather than by acting directly.

2.3 Predictive Modeling and Learning Analytics

Educational data mining and learning analytics have established predictive modeling as a practical tool for anticipating student outcomes. Comparative studies benchmark algorithms ranging from logistic and linear regression to support vector machines, random forests, gradient boosting, and neural networks, typically reporting that ensemble and, in some datasets, well-specified linear models offer the best balance of accuracy and stability (Yağcı, 2022; Ahmed et al., 2025). A recurring finding is that behavioral and self-regulatory indicators—engagement consistency, assignment timeliness, prior achievement—dominate demographic variables in predictive power (Ahmed et al., 2025).

The interpretability of these models has become a first-order concern. Techniques such as SHAP and permutation importance allow institutions to understand why a prediction is made, transforming opaque classifiers into decision-support instruments; recent work shows that a small subset of variables can explain the majority of performance variance, and that policy simulations built on interpretable models can quantify the expected impact of interventions (Guevara-Reyes et al., 2025). Complementary research argues that the true value of AI prediction is realized only when it is coupled with learning analytics feedback that informs pedagogical action, closing the

loop between forecasting and instruction (Ouyang et al., 2023). The present study adopts this interpretable, action-oriented stance by pairing a benchmarking exercise with permutation-based feature attribution and an explicit test of moderation.

2.4 Hypotheses

Integrating the strands above, five hypotheses are advanced. Consistent with meta-analytic evidence of a positive but modest aggregate effect (Sun & Zhou, 2024), and with survey findings of weak main effects (Aslam et al., 2025):

H1. AI use intensity is positively but weakly associated with the quality of learning at the zero-order level.

H2. Self-regulated learning is a stronger predictor of learning quality than AI use intensity or AI literacy (Shi et al., 2025).

H3. Engagement is among the strongest proximal predictors of learning quality, consistent with its mediating role (Hassan, 2026).

H4. Pedagogical guidance moderates the AI-use–learning-quality relationship, such that AI use predicts higher quality under structured use than under unguided use (Hassan, 2026).

H5. Interpretable and ensemble machine-learning models achieve comparable, useful predictive accuracy for learning quality, with a small set of self-regulatory and behavioral features dominating importance (Ahmed et al., 2025; Guevara-Reyes et al., 2025).

3. Materials and Methods

3.1 Research Design

The study employs a quantitative, non-experimental, cross-sectional, and predictive design, following the projective-predictive tradition used in recent AI-and-performance research in higher education (Pacheco-Mendoza et al., 2023). The design is explanatory-predictive: inferential regression is used to test the hypothesized relationships and moderation, while a

supervised machine-learning benchmark evaluates out-of-sample predictive performance and feature importance. This dual approach mirrors the methodological consensus that combines interpretable inference with predictive validation (Guevara-Reyes et al., 2025).

3.2 Data and Sample

Transparency statement. Because primary data collection across multiple institutions was outside the scope of this methodological study, the analysis was conducted on a simulated cohort of 1,200 undergraduate students. The data-generating process was calibrated to relationships, directions, and effect sizes reported in the peer-reviewed literature cited throughout Section 2 (e.g., the medium aggregate AI effect of Sun & Zhou, 2024; the SRL-over-AI-literacy ordering of Shi et al., 2025; and the guidance-dependent effect of Hassan, 2026). The simulated data are used strictly to demonstrate and stress-test the predictive-modeling pipeline; the reported coefficients illustrate the framework rather than estimate population parameters for any specific institution. All code, parameters, and the generated dataset are available (see Data Availability) so that the pipeline can be re-run on real institutional data.

The simulated cohort approximates a typical undergraduate population: mean age 21.6 years (SD = 3.0), a balanced distribution of pedagogical-guidance conditions (43% structured use), and continuous distributions for AI use, literacy, self-regulation, engagement, study hours, and socioeconomic status. The sample size exceeds the minimum required for stable estimation of the nine-predictor models and supports 10-fold cross-validation with adequate fold sizes.

3.3 Variables and Measures

The outcome variable, quality of learning, is a composite scaled 0–100 that integrates mastery, engagement, and cognitive investment. Nine predictors span AI-related, self-regulatory, behavioral, and background domains. Table 1 summarizes the operationalization of each variable and its theoretical anchor.

Table 1. Operationalization of study variables.

Variable	Operationalization (scale)	Theoretical anchor
Quality of learning	Composite of mastery, engagement, cognitive investment (0–100)	Hassan (2026); Anghel et al. (2025)
AI use intensity	Self-reported hours/week using AI learning tools	Sun & Zhou (2024)

Variable	Operationalization (scale)	Theoretical anchor
AI literacy	Competence in using/evaluating AI tools (1–5)	Shi et al. (2025)
Self-regulated learning	Planning, monitoring, adjusting learning (1–5)	Zimmerman (2002); Shi et al. (2025)
Pedagogical guidance	Structured vs. unguided AI use (0/1)	Hassan (2026)
Engagement index	Platform activity / participation (0–100)	Ouyang et al. (2023)
Prior GPA	Prior academic achievement (0–20)	Ahmed et al. (2025)
Study hours	Weekly study time (hours)	Pacheco-Mendoza et al. (2023)
Socioeconomic index	Composite SES indicator (1–5)	Guevara-Reyes et al. (2025)
Age	Student age (years)	Pacheco-Mendoza et al. (2023)

Note. Scales reflect the operationalization used in the simulated pipeline; on real data, validated instruments (e.g., SRL and AI-literacy scales) would replace the corresponding composites.

3.4 Predictive Models

Five regression algorithms were benchmarked, spanning the interpretability–flexibility spectrum: (i) multiple linear regression (MLR) as an interpretable baseline; (ii) random forest (RF); (iii) gradient boosting (GB); (iv) extreme gradient boosting (XGBoost); and (v) an artificial neural network (ANN) implemented as a multilayer perceptron with two hidden layers (64 and 32 units, ReLU activation, Adam optimizer). This set covers the algorithms most frequently reported in educational prediction studies (Yağcı, 2022; Ahmed et al., 2025). Data were partitioned into 75% training and 25% hold-out test sets. Continuous features were standardized for the ANN; tree-based models used raw features. Hyperparameters were fixed at conventional, lightly tuned values to avoid overfitting the demonstration set.

3.5 Evaluation Protocol

Model performance was assessed on the hold-out test set using the coefficient of determination (R^2), root-mean-square error (RMSE), and mean absolute error (MAE), complemented by 10-fold cross-validation (R^2 mean and SD) on the full dataset to gauge stability. Permutation importance (30 repeats) was computed on the XGBoost model to attribute predictive contribution to each feature in a model-agnostic manner, consistent with the interpretability emphasis of contemporary work (Guevara-Reyes et al., 2025).

Moderation (H4) was tested with an ordinary-least-squares model including an AI-use $\times$ pedagogical-guidance interaction term, controlling for all remaining predictors. Statistical significance was evaluated at $\alpha = 0.05$.

3.6 Ethical Considerations

Although the analysis used simulated data, the framework is designed for eventual deployment on real student records, which raises well-documented ethical obligations: informed consent, data privacy, transparency, and the mitigation of algorithmic bias (Remegio, 2025; Msambwa et al., 2025). Predictive early-warning systems must be used to support—not stigmatize—students, with human oversight of any high-stakes decision. Interpretable models are preferred precisely because they permit scrutiny of the factors driving a prediction, reducing the risk of opaque or discriminatory inferences.

4. Results

4.1 Descriptive Statistics and Correlations

Table 2 reports descriptive statistics for the outcome and predictors. The quality-of-learning composite averaged 68.6 (SD = 9.5) on the 0–100 scale, with the full range represented. AI use averaged 4.6 hours per week (SD = 3.1), and 43% of the cohort experienced structured, guided AI use.

Table 2. Descriptive statistics (N = 1,200).

Variable	Mean	SD	Min	Max
Quality of learning (0–100)	68.57	9.48	37.65	100.00
Prior GPA (0–20)	14.06	2.29	6.15	20.00
AI use (hours/week)	4.60	3.10	0.15	18.04
AI literacy (1–5)	3.23	0.84	1.00	5.00
Self-regulated learning (1–5)	3.35	0.77	1.00	5.00
Pedagogical guidance (0/1)	0.43	0.50	0.00	1.00
Engagement index (0–100)	49.16	10.10	14.20	86.98
Study hours (per week)	13.19	4.65	0.00	30.62
Age (years)	21.60	2.98	17.00	32.00
Socioeconomic index (1–5)	2.97	0.88	1.00	5.00

Zero-order Pearson correlations with learning quality (Figure 1) were consistent with the hypotheses. Engagement ($r = 0.581$, $p < 0.001$) and self-regulated learning ($r = 0.544$, $p < 0.001$) showed the strongest associations, followed by AI literacy ($r = 0.322$, $p < 0.001$) and prior GPA ($r = 0.258$, $p < 0.001$). AI use

intensity showed a positive but comparatively weak correlation ($r = 0.164$, $p < 0.001$), supporting H1. Pedagogical guidance ($r = 0.202$, $p < 0.001$) and study hours ($r = 0.224$, $p < 0.001$) were modest positive correlates, while age was not significantly related to quality ($r = -0.013$, $p = 0.65$).




Figure 1. Pearson correlation matrix of the outcome and nine predictors (N = 1,200).

4.2 Inferential Model: Ordinary Least Squares

The OLS model (Table 3) explained a substantial share of variance and clarified the partial contributions of each predictor. Self-regulated learning had the largest standardized influence ($b = 4.66$, $p < 0.001$), followed by pedagogical guidance ($b = 2.42$, $p < 0.001$), AI literacy ($b = 2.14$, $p < 0.001$), prior GPA ($b = 1.09$, $p < 0.001$), and socioeconomic index ($b = 1.03$, $p <$

0.001). Engagement ($b = 0.27$, $p < 0.001$) and study hours ($b = 0.38$, $p < 0.001$) were significant per-unit contributors given their wider scales. Notably, AI use intensity retained a small but significant partial coefficient ($b = 0.24$, $p < 0.001$), whereas age was non-significant ($p = 0.19$). This pattern supports H1–H3: AI use matters, but its unique contribution is dwarfed by self-regulation and engagement.

Table 3. Multiple linear regression predicting quality of learning.

Predictor	Coef.	Std. err.	t	p-value
Constant	8.999	2.141	4.203	< 0.001
Prior GPA	1.088	0.077	14.115	< 0.001
AI use (hours/week)	0.237	0.059	4.002	< 0.001
AI literacy	2.137	0.221	9.678	< 0.001
Self-regulated learning	4.663	0.268	17.392	< 0.001
Pedagogical guidance	2.415	0.373	6.469	< 0.001
Engagement index	0.269	0.022	12.078	< 0.001
Study hours	0.379	0.039	9.847	< 0.001
Age	-0.078	0.059	-1.313	0.190
Socioeconomic index	1.027	0.201	5.116	< 0.001

Note. Dependent variable: quality of learning (0–100). $N = 1,200$.

4.3 Predictive Model Benchmark

Table 4 and Figure 2 compare the five algorithms. Multiple linear regression achieved the best hold-out performance ($R^2 = 0.567$; RMSE = 6.04; MAE = 4.75) and the highest cross-validated R^2 (0.579 ± 0.059), narrowly outperforming gradient boosting ($R^2 = 0.524$), random forest ($R^2 = 0.516$), and XGBoost ($R^2 = 0.502$). The ANN underperformed ($R^2 = 0.351$) and

showed the greatest cross-validation variance (SD = 0.109), a common outcome for neural networks on tabular data of this size. The strong showing of the linear model echoes recent findings that well-specified linear regression can rival or exceed ensemble methods on structured educational data (Ahmed et al., 2025), and it supports H5: interpretable models delivered useful, stable accuracy.

Table 4. Predictive performance of candidate models.

Model	R^2 (test)	RMSE	MAE	10-fold CV R^2
Multiple linear regression	0.567	6.038	4.748	0.579 ± 0.059
Gradient boosting	0.524	6.326	4.832	0.526 ± 0.061
Random forest	0.516	6.383	4.999	0.523 ± 0.047
XGBoost	0.502	6.472	5.043	0.509 ± 0.066
Artificial neural network (MLP)	0.351	7.389	5.905	0.417 ± 0.109




Figure 2. Hold-out R^2 (bars) and 10-fold cross-validated R^2 (points, mean $\pm$ SD) by model.

The observed-versus-predicted plot for the best model (Figure 4) shows values distributed closely around the identity line without systematic bias,

indicating well-calibrated predictions across the range of learning quality.




Figure 4. Observed versus predicted quality of learning on the hold-out test set (multiple linear regression).

4.4 Feature Importance

Permutation importance from the XGBoost model (Table 5, Figure 3) provides a model-agnostic view of predictive contribution. Three predictors dominate: self-regulated learning (32.2%), engagement (21.3%), and prior GPA (20.1%), which together account for roughly three-quarters of total importance. Study hours (10.3%) and AI literacy (9.2%) form a second

tier. Strikingly, the two most direct AI-usage variables—raw AI use intensity (2.5%) and pedagogical guidance considered as a main effect (3.1%)—contribute modestly, while age is negligible. This ordering strongly supports H2 and H3 and refines H1: at the level of main effects, how much AI a student uses matters far less than whether that student self-regulates and engages.

Table 5. Permutation feature importance (XGBoost, 30 repeats).

Feature	Importance	Relative (%)
Self-regulated learning	0.244	32.16
Engagement index	0.162	21.34
Prior GPA	0.153	20.09
Study hours	0.078	10.30
AI literacy	0.070	9.22
Pedagogical guidance	0.024	3.09
AI use (hours/week)	0.019	2.46
Socioeconomic index	0.011	1.42
Age	-0.001	-0.09




Figure 3. Relative permutation importance of predictors; AI-related predictors highlighted.

4.5 Moderation: AI Use × Pedagogical Guidance

The apparently modest main effect of AI use masks a meaningful conditional relationship. The OLS interaction model returned a positive and significant AI-use × pedagogical-guidance term ($b = 0.447$, $p < 0.001$), confirming H4. As Figure 5 illustrates, additional hours of AI use are associated with rising learning quality under structured, guided conditions,

but with a flat or slightly declining trajectory under unguided use. In other words, the same quantity of AI use yields divergent outcomes depending on pedagogical framing: AI amplifies learning when scaffolded and adds little—or detracts—when left to unsupported self-direction. This finding reconciles the meta-analytic optimism of Sun and Zhou (2024) with the cautionary evidence of Hassan (2026) and the weak main effects reported by Aslam et al. (2025).




Figure 5. Mean quality of learning across AI-use intensity, by pedagogical-guidance condition.

4.6 Summary of Hypothesis Tests

Table 6. Summary of hypothesis testing.

H	Statement (abbreviated)	Outcome
H1	AI use is positively but weakly associated with learning quality	Supported
H2	SRL is a stronger predictor than AI use or AI literacy	Supported
H3	Engagement is among the strongest proximal predictors	Supported
H4	Pedagogical guidance moderates the AI-use–quality relationship	Supported
H5	Interpretable/ensemble models give comparable, useful accuracy	Supported

Note. “Supported” denotes consistency of the estimated pattern with the hypothesis within the demonstration pipeline.

5. Discussion

5.1 Principal Findings

This study set out to characterize the relationship between AI use and the quality of learning in higher education and to identify, via predictive modeling, the factors that most strongly forecast learning quality. Three findings stand out. First, AI use intensity is, on its own, only a weak predictor of learning quality: its zero-order correlation was modest ($r = 0.164$) and its permutation importance was among the lowest (2.5%). Second, the strongest predictors are learner-level self-regulatory and engagement variables—self-regulated learning and engagement jointly accounted for more than half of predictive importance—alongside prior achievement. Third, and most consequentially, the effect of AI use is conditional: a significant interaction with pedagogical guidance showed that AI use predicts higher learning quality

only when embedded in structured pedagogical support.

Taken together, these results support a reframing of AI's role. Rather than acting as an autonomous driver of learning quality, AI functions as an amplifier of pre-existing learning conditions. Where self-regulation, engagement, and scaffolding are present, additional AI use is associated with better learning; where they are absent, more AI use does little to help and may even substitute for the effortful processing that produces durable learning.

5.2 Theoretical Implications

The findings extend acceptance- and outcome-oriented theories of educational technology by foregrounding moderation. Models grounded in the Unified Theory of Acceptance and Use of Technology explain adoption well (Dahri et al., 2024; Venkatesh et al., 2003), but adoption is not synonymous with

beneficial learning. Our results suggest that theories of AI in education must incorporate the pedagogical and self-regulatory conditions under which use translates into quality, echoing evidence that self-regulated learning predicts performance more strongly than AI literacy (Shi et al., 2025) and that engagement mediates the AI–achievement link (Hassan, 2026).

The result also speaks to the construct of learning quality itself. When quality is operationalized to include engagement and cognitive investment—not merely grades—AI's contribution appears indirect, channeled through the conditions that foster deep processing. This aligns with dual-edged accounts in which AI simultaneously stimulates and constrains higher-order cognition depending on how it is used (Lin et al., 2024), and with the observation that unguided reliance can undermine the very problem-solving it appears to assist (Hassan, 2026).

5.3 Practical Implications for Higher Education Institutions

For institutional leaders, the message is not to maximize AI adoption per se but to invest in the conditions that make AI beneficial. Three levers follow directly from the results. First, integrate AI within structured pedagogy: the significant interaction indicates that scaffolding, guided tasks, and instructor framing convert AI use into learning gains, consistent with calls for instructional scaffolding around generative tools (Hassan, 2026). Second, develop self-regulation and AI literacy jointly, since these learner capacities—especially self-regulation—are the strongest malleable predictors of quality (Shi et al., 2025). Third, deploy predictive early-warning systems built on interpretable models: because a small set of self-regulatory and behavioral indicators explains most of the variance, institutions can monitor these signals to identify at-risk students early and target support, closing the loop between prediction and pedagogical action (Ouyang et al., 2023; Guevara-Reyes et al., 2025).

The competitive performance of the interpretable linear model is itself practically valuable. Institutions frequently face a trade-off between accuracy and explainability; here, the most transparent model was also the most accurate and stable, lowering the barrier to responsible deployment and enabling straightforward communication of predictions to advisors and students (Remegio, 2025).

5.4 Comparison with Prior Work

The results occupy a coherent middle ground in a divided literature. They are consistent with meta-

analytic evidence that AI can improve achievement (Sun & Zhou, 2024) while explaining why some primary studies find weak or null main effects (Aslam et al., 2025): the aggregate effect is real but conditional, and averaging over guided and unguided contexts attenuates it. The dominance of self-regulatory and engagement features replicates the feature-importance patterns reported in explainable performance-prediction studies, where behavioral and prior-achievement indicators outrank demographics (Ahmed et al., 2025). Finally, the predictive framework operationalizes the integration of AI prediction with learning analytics advocated by Ouyang et al. (2023), demonstrating that interpretable models can serve both explanation and forecasting.

6. Limitations and Future Research

Several limitations qualify these findings and define an agenda for future work. The most important is that the analysis used a simulated cohort calibrated to the literature; although this permits transparent demonstration of the pipeline and controlled testing of moderation, the specific coefficients should not be read as population estimates. The immediate priority is replication on multi-institutional primary data using validated instruments for self-regulated learning, AI literacy, and engagement, together with objective platform-derived behavioral traces rather than self-reports.

Second, the cross-sectional design precludes causal inference; longitudinal and quasi-experimental designs that compare pre-AI, unguided-AI, and structured-AI phases (Hassan, 2026) would strengthen causal claims and capture the temporal dynamics of learning. Third, the outcome was a composite; future work should disaggregate quality into distinct facets—critical thinking, transfer, creativity, and well-being—since AI's effects appear to differ across them (Lin et al., 2024; Shi et al., 2025). Fourth, richer models could test mediation formally (AI use → engagement → quality) and additional moderators such as discipline, task type, and institutional readiness (Aslam et al., 2025). Finally, deployment on real data must be accompanied by rigorous attention to privacy, consent, fairness auditing, and human oversight (Remegio, 2025; Msambwa et al., 2025).

7. Conclusion

The relationship between AI use and the quality of learning in higher education is best understood as conditional rather than direct. Across an interpretable predictive framework, AI use intensity emerged as a weak stand-alone predictor of learning quality, while

self-regulated learning, engagement, and prior achievement dominated—and, critically, the benefit of AI use materialized only under structured pedagogical guidance. Interpretable and ensemble models delivered comparable, useful, and stable accuracy, with the most transparent model performing best, underscoring that explainability and predictive value need not be traded off in educational analytics.

For higher education institutions, the practical implication is clear: the returns to AI depend on the learning conditions surrounding it. Investing in students' self-regulation, embedding AI within scaffolded pedagogy, and adopting interpretable early-warning systems will do more to raise learning quality than expanding AI access alone. Framed this way, AI is neither a panacea nor a threat but an amplifier—one whose value HEIs can deliberately cultivate through sound pedagogy and evidence-based decision-making.

References

- Adewale, M. D., Azeta, A., Abayomi-Alli, A., & Sambo-Magaji, A. (2024). Impact of artificial intelligence adoption on students' academic performance in open and distance learning: A systematic literature review. *Heliyon*, 10(22), e40025.
- Ahmed, W., Alhassan, S., & colleagues. (2025). Machine learning-based academic performance prediction with explainability for enhanced decision-making in educational institutions. *Scientific Reports*, 15.
- Anghel, G. A., Ariton, A., & colleagues. (2025). The integration of artificial intelligence in academic learning practices: A comprehensive approach. *Education Sciences*, 15.
- Aslam, F., & colleagues. (2025). Artificial intelligence in higher education: Shaping the future of university teaching through adaptive learning, intelligent tutoring, and academic analytics. *Indus Journal of Social Sciences*, 3.
- Chaudhry, I. S., Sarwary, S. A. M., El Refae, G. A., & Chabchoub, H. (2023). Time to revisit existing student's performance evaluation approach in higher education sector in a new era of ChatGPT — A case study. *Cogent Education*, 10(1).
- Chen, L., Chen, P., & Lin, Z. (2020). Artificial intelligence in education: A review. *IEEE Access*, 8, 75264–75278. <https://doi.org/10.1109/ACCESS.2020.2988510>
- Crompton, H., & Burke, D. (2023). Artificial intelligence in higher education: The state of the field. *International Journal of Educational Technology in Higher Education*, 20(1), 22. <https://doi.org/10.1186/s41376-1388-01376-1388>
- Dahri, N. A., & colleagues. (2024). Investigating AI-based academic support acceptance and its impact on students' performance in Malaysian and Pakistani higher education institutions. *Education and Information Technologies*, 29.
- Guevara-Reyes, R., Ortiz-Garcés, I., & Andrade, R. (2025). Machine learning models for academic performance prediction: Interpretability and application in educational decision-making. *Frontiers in Education*, 10, 1632315. <https://doi.org/10.3389/feduc.2025.1632315>
- Hassan, A. (2026). Leveraging artificial intelligence (AI) to enhance student engagement and academic performance in higher education. *Education and Information Technologies*. Advance online publication.
- Lin, H., & colleagues. (2024). Artificial intelligence (AI)-integrated educational applications and college students' creativity and academic emotions: Students and teachers' perceptions and attitudes. *BMC Psychology*, 12.

Declarations

Funding. This research received no external funding.

Institutional Review Board Statement. Not applicable; the study used simulated data and did not involve human participants.

Informed Consent Statement. Not applicable.

Data Availability Statement. The simulated dataset, the data-generating parameters, and all analysis and figure code are available from the corresponding author on reasonable request, enabling replication of the pipeline on real institutional data.

Conflicts of Interest. The authors declare no conflict of interest.

Note on citation verification. All cited empirical sources were checked against the Scite and Consensus scholarly databases; DOIs for the central sources were confirmed and are listed in the References.

- Msambwa, M. M., Wen, Z., & Daniel, K. (2025). The impact of AI on the personal and collaborative learning environments in higher education. *European Journal of Education*, 60.
- Ouyang, F., Wu, M., Zheng, L., Zhang, L., & Jiao, P. (2023). Integration of artificial intelligence performance prediction and learning analytics to improve student learning in online engineering course. *International Journal of Educational Technology in Higher Education*, 20(1), 4. <https://doi.org/10.1186/s41376-1388-01376-1388>
- Pacheco-Mendoza, S., Guevara, C., & Mayorga Albán, A. L. (2023). Artificial intelligence in higher education: A predictive model for academic performance. *Education Sciences*, 13(10), 990. <https://doi.org/10.3390/educsci13100990>
- Remegio, A. N. (2025). Predictive analytics in higher education: A systematic literature review of machine learning models for forecasting academic performance and ethical integration. In *2025 14th International Conference on Educational and Information Technology (ICEIT)*. IEEE.
- Shi, J., & colleagues. (2025). Exploring how AI literacy and self-regulated learning relate to student writing performance and well-being in generative AI-supported higher education. *Behavioral Sciences*, 15.
- Slimi, Z. (2023). The impact of artificial intelligence on higher education: An empirical study. *European Journal of Educational Sciences*, 10(1), 17–33.
- Sun, L., & Zhou, L. (2024). Does generative artificial intelligence improve the academic achievement of college students? A meta-analysis. *Journal of Educational Computing Research*, 62(7), 1676–1713. <https://doi.org/10.1177/07356331241277937>
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425–478.
- Yağcı, M. (2022). Educational data mining: Prediction of students' academic performance using machine learning algorithms. *Smart Learning Environments*, 9(1), 11.
- Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education — Where are the educators? *International Journal of Educational Technology in Higher Education*, 16(1), 39. <https://doi.org/10.1186/s41376-1388-1376-1388>
- Zimmerman, B. J. (2002). Becoming a self-regulated learner: An overview. *Theory Into Practice*, 41(2), 64–70.