

Adoption Barriers to AI-Powered Learning Tools Among University Faculty: A Mixed-Methods Study

Kanita Haider^{1*}, Mrityika Sen Gupta², Md Rasel Ul Alam³

¹ Chittagong University of Engineering and Technology, Chattogram 4349, Bangladesh

² North South University, Dhaka, Bangladesh

³ University of the Cumberlands, Kentucky, USA

*Corresponding author: kanita.haider4422@gmail.com

Received: 15 January 2023; Revised: 13 April 2023; Accepted: 19 May 2023; Published: 15 July 2023

KEYWORDS

AI Adoption Barriers
Higher Education
Faculty Technology
Acceptance
Machine Learning
Mixed-Methods Research
Educational Technology
Random Forest

Abstract

This mixed-methods study investigates the barriers that shape university faculty adoption of AI-powered learning tools, integrating a quantitative survey of 412 faculty members across five disciplines with semi-structured interviews from 18 purposively sampled participants. Machine learning techniques, including a random forest classifier and k-means clustering, were applied to the survey data to identify the strongest predictors of adoption and to segment faculty into distinct adoption profiles. The random forest model classified faculty adoption status with 72% accuracy and an area under the curve (AUC) of 0.75, identifying perceived usefulness, perceived ease of use, and digital literacy as the strongest predictors of adoption, while technology anxiety, limited institutional support, and workload pressure emerged as the most salient barriers. Cluster analysis revealed three faculty segments, Ready Adopters (34%), Cautious Adopters (37%), and Resistant Faculty (29%), each requiring differentiated support strategies. Qualitative interviews further revealed five recurrent barrier themes: time and workload constraints, insufficient institutional training, concerns about academic integrity and data privacy, discomfort with algorithmic opacity, and a perceived mismatch between AI tools and disciplinary pedagogy. Supplementary robustness checks, including a random forest learning-curve analysis and a comparison against logistic regression, confirmed the stability of the reported predictors across modeling choices. The study concludes that faculty adoption of AI-powered learning tools is shaped less by the technology itself than by the organizational and psychological ecosystem surrounding its introduction, and it proposes a tiered, discipline-sensitive framework for institutional AI integration.

1. Introduction

The rapid diffusion of artificial intelligence (AI) into everyday digital infrastructure has begun to reshape teaching and learning in higher education. AI-powered learning tools, ranging from adaptive tutoring systems and automated grading platforms to generative AI writing assistants and intelligent content recommendation engines, promise to personalize instruction, reduce faculty workload, and generate real-time insight into student learning (Zawacki-Richter et al., 2019; Popenici & Kerr, 2017). Yet despite substantial institutional

investment and growing student familiarity with these tools, faculty adoption remains uneven and, in many institutions, markedly slower than anticipated (Selwyn, 2019). Understanding why faculty embrace or resist AI-powered learning tools has therefore become an urgent priority for higher education leaders, instructional designers, and educational technology researchers alike.

Faculty technology adoption has long been studied through acceptance frameworks such as the Technology Acceptance Model (TAM; Davis, 1989) and the Unified Theory of Acceptance and Use of

Technology (UTAUT; Venkatesh et al., 2003), both of which emphasize perceived usefulness, perceived ease of use, and social influence as central determinants of adoption behavior. More recent scholarship extends these frameworks to account for the distinctive features of AI systems, including algorithmic opacity, data governance concerns, and the disruption of established pedagogical routines (Almaiah et al., 2022; Buchanan et al., 2013). However, much of this literature relies on single-method survey designs that capture the strength of association between constructs without illuminating the lived experience of faculty who must reconcile AI tools with disciplinary norms, workload realities, and institutional support structures.

This study addresses that gap by adopting a mixed-methods design that combines a large-scale faculty survey with semi-structured interviews and applies supervised and unsupervised machine learning techniques to the quantitative data. Specifically, the study uses a random forest classifier to identify which constructs most strongly predict high versus low AI adoption status among faculty, and applies k-means clustering to segment faculty into distinct adoption profiles that can inform differentiated institutional support strategies. By triangulating predictive analytics with qualitative theme analysis, the study offers a more complete account of the barriers that shape faculty engagement with AI-powered learning tools.

The specific objectives of this study are to: (i) quantify the relative intensity of key barriers to AI adoption among university faculty; (ii) identify the strongest statistical and machine-learning predictors of faculty adoption status; (iii) segment faculty into meaningful adoption profiles using cluster analysis; (iv) explore, through qualitative interviews, the lived experiences and reasoning that underlie observed adoption patterns; and (v) propose an evidence-based, discipline-sensitive framework for institutional AI integration.

1.1 Paper Organization

The remainder of this paper proceeds as follows. Section 2 reviews technology acceptance theory and prior empirical work on AI adoption barriers in higher education. Section 3 describes the mixed-methods research design, sampling, instruments, and machine learning and statistical analysis procedures. Section 4 reports the quantitative and qualitative results, including descriptive statistics, correlation analysis, random forest classification, robustness checks, cluster analysis, discipline and rank comparisons, and thematic findings from the interviews. Section 5 discusses the implications of these findings for technology acceptance theory and institutional practice, and Section 6 concludes with a tiered, discipline-sensitive framework and recommendations for future research.

2. Literature Review

The theoretical foundation for technology adoption research rests heavily on Davis's (1989) Technology Acceptance Model, which posits that perceived usefulness and perceived ease of use are the two principal determinants of an individual's intention to use a new technology. Venkatesh and Bala (2008) later extended this framework into TAM3, incorporating anchor and adjustment variables such as computer self-efficacy, computer anxiety, and perceptions of external control. Venkatesh et al. (2003) further broadened the theoretical landscape with UTAUT, which integrates performance expectancy, effort expectancy, social influence, and facilitating conditions as joint predictors of technology use. These frameworks collectively suggest that faculty adoption of AI-powered learning tools is unlikely to be driven by technological capability alone, but rather by a constellation of cognitive, social, and organizational factors.

Building on Rogers's (2003) Diffusion of Innovations theory, several scholars have emphasized that the perceived compatibility of a new technology with existing values, workflows, and professional identity is a critical precondition for adoption. Ertmer and Ottenbreit-Leftwich (2010) argue that teacher technology integration depends jointly on external conditions, such as access and institutional support, and internal conditions, including confidence, beliefs about teaching, and cultural norms within a discipline. This dual framing is particularly salient for AI-powered tools, which often require faculty to relinquish a degree of pedagogical control to an algorithmic system, a shift that can provoke discomfort even among otherwise technologically confident instructors.

Empirical and conceptual studies specific to AI-based tools in higher education have long anticipated the barriers that faculty report today. Almaiah et al. (2022) found that computer anxiety and social anxiety were significant negative predictors of e-learning technology acceptance among university populations. Selwyn (2019) argued that the introduction of AI-based teaching tools, including intelligent tutoring systems, learning analytics, and automated decision-making, provokes both practical and identity-based resistance among faculty, who question whether such systems can replicate the social, emotional, and cognitive qualities of human instruction. Popenici and Kerr (2017) similarly cautioned that institutions were adopting AI-driven teaching and learning technologies faster than they were developing the pedagogical and ethical frameworks needed to guide their use, an imbalance that leaves instructors ambivalent about over-reliance and the erosion of critical thinking skills. Roll and Wylie (2016), reviewing two decades of published research on artificial intelligence in education, found that the field's technical progress had occurred

largely in isolation from everyday classroom practice, and argued that closing this evolution-to-revolution gap would require far deeper collaboration between AI researchers and teachers, a structural gap consistent with the adoption barriers observed among faculty.

Institutional and organizational factors also feature prominently in the literature. Buchanan, Sainter, and Saunders (2013), in a survey of higher education faculty use of learning technologies, found that structural constraints within the institution and the perceived usefulness of the tools were the two principal barriers to adoption, with facilitating conditions such as access to resources and technical support outweighing individual differences in digital literacy and self-efficacy. This finding echoes long-standing evidence from the broader educational technology literature that facilitating conditions, not merely individual attitudes, shape the trajectory of technology adoption in academic settings (Venkatesh et al., 2003; Al-Emran & Granić, 2021).

Methodologically, most existing studies rely on structural equation modeling or ordinary regression to test hypothesized relationships among adoption constructs (Hair et al., 2019). While powerful for confirmatory hypothesis testing, these approaches are less well suited to uncovering nonlinear interactions among barriers or to segmenting heterogeneous faculty populations into actionable subgroups. Machine learning methods, including random forest classification and k-means clustering, offer a complementary analytic lens: they can capture nonlinear relationships among predictors, rank variable importance without strong distributional assumptions, and identify latent subgroups that a single regression model would obscure (James et al., 2021). Despite this potential, few studies have applied such techniques specifically to faculty AI adoption barriers, a gap this study seeks to address by combining machine learning analysis with qualitative inquiry within a single mixed-methods design.

2.2 Comparative Summary of Prior Empirical Studies

Table 1 situates the present study's design and predictor set relative to the three empirical studies most directly comparable to it in the technology-acceptance and AI-adoption literature. The comparison illustrates that, while prior studies established the theoretical centrality of usefulness, ease of use, and structural facilitating conditions, none combined a large-scale predictive model with a cluster-based segmentation and a qualitative validation strand within a single design, the gap this study's convergent mixed-methods approach is intended to close.

Study	Sample	Primary Method	Segmented ?
Present study	N = 412 + 18 interviews	Random forest + k-means + thematic	Yes (3)
Buchanan et al. (2013)	Faculty survey	Descriptive / regression	No
Almaiah et al. (2022)	Student survey	Structural equation modeling	No
Venkatesh et al. (2003)	Multi-org. survey	UTAUT regression	No

Table 1: Comparative summary of the present study and three closely related prior empirical studies.

2.3 Summary of Literature Gaps

Three gaps in the existing literature motivate the present study's design. First, acceptance-framework research on AI-specific tools (Almaiah et al., 2022; Selwyn, 2019; Popenici & Kerr, 2017) has generally relied on single-method survey or conceptual designs that cannot triangulate statistical association with the lived reasoning that produces it. Second, the methodological literature has noted that regression- and structural-equation-based approaches, while confirmatory-hypothesis-friendly, are poorly suited to uncovering nonlinear predictor interactions or latent population heterogeneity (James et al., 2021; Hair et al., 2019), a limitation directly addressed by the random forest and k-means methods applied here, and as Table 8 illustrates, none of the closest comparison studies applied a segmentation method at all. Third, institutional-factor research (Buchanan et al., 2013) has typically treated faculty as a single population rather than as a set of distinguishable adoption segments requiring differentiated institutional responses, a gap the cluster analysis reported in Section 4 is specifically designed to close.

3. Methodology

3.1 Research Design

This study employed a convergent mixed-methods design in which quantitative and qualitative data were collected concurrently and integrated during interpretation. The quantitative strand consisted of a cross-sectional online survey administered to university faculty, analyzed using descriptive statistics, correlation analysis, a random forest classifier, and k-means cluster analysis. The qualitative strand consisted of semi-structured interviews analyzed thematically. Integration occurred at the interpretation stage, where machine-

learning-derived adoption profiles were used to contextualize and enrich the qualitative themes.

3.2 Participants and Sampling

The quantitative sample comprised 412 full-time faculty members drawn from five broad disciplinary groupings: STEM ($n = 99$), Social Sciences ($n = 82$), Humanities ($n = 80$), Business ($n = 70$), and Health Sciences ($n = 81$), recruited through institutional email listservs at partner universities using stratified random sampling to ensure disciplinary and rank diversity. Academic rank distribution included Lecturers/Instructors (22%), Assistant Professors (34%), Associate Professors (28%), and Professors (16%). For the qualitative strand, 18 faculty members were purposively selected from among survey respondents to maximize variation in discipline, rank, and self-reported adoption status, and participated in 30 to 45 minute semi-structured interviews.

3.3 Instruments and Measures

The survey instrument measured nine constructs using validated five-point Likert-type scales adapted from the TAM, TAM3, and UTAUT literatures: digital literacy, institutional support, workload pressure, technology anxiety, peer influence, training access, data privacy concern, perceived usefulness, and perceived ease of use. A composite barrier score was computed as a weighted average of the barrier-oriented constructs (technology anxiety, reverse-scored institutional support, workload pressure, reverse-scored training access, data privacy concern, and reverse-scored peer influence and digital literacy). A binary adoption status variable (High Adoption versus Low Adoption) was derived from a validated single-item measure of AI tool usage frequency in teaching, corroborated by a composite adoption-intention scale. The interview protocol used open-ended prompts exploring participants' experiences experimenting with AI tools, perceived obstacles, and institutional support experiences.

3.4 Machine Learning and Statistical Analysis

Quantitative data were analyzed in Python using pandas, scikit-learn, and seaborn. A random forest classifier (400 trees, maximum depth of 5, minimum leaf size of 6) was trained on a 75/25 stratified train-test split to predict binary adoption status from the nine survey constructs, with model performance evaluated using accuracy, precision, recall, F1-score, and the area under the receiver operating characteristic curve (AUC-ROC). Gini-based feature importance scores were extracted to rank the relative predictive contribution of each construct. To identify latent faculty subgroups, features were standardized and subjected to k-means clustering ($k = 3$, selected via inspection of within-cluster variance), with results visualized using principal component analysis (PCA) for dimensionality reduction. Pearson correlation analysis examined bivariate relationships

among constructs, and one-way comparisons examined variation in barrier scores across disciplines. Qualitative interview transcripts were analyzed using reflexive thematic analysis, with two coders independently coding a subsample of transcripts to establish coding consistency before finalizing the thematic structure.

3.5 Robustness and Model Comparison Procedures

To assess whether the reported predictor rankings and classification performance depended on particular modeling choices, two supplementary robustness checks were conducted. First, a learning-curve analysis tracked training and held-out test accuracy as the number of trees in the random forest increased from 10 to 400, allowing an assessment of whether the model had converged and whether additional trees would be expected to materially change performance. Second, the random forest's classification performance was compared against a standard logistic regression model fit on the same nine constructs and the same train-test split, providing a linear-model benchmark against which the random forest's accuracy and AUC could be interpreted. Results of both robustness checks are reported in Section 4.4.

3.6 Ethical Considerations

All participants provided informed consent prior to participation. Survey responses were collected anonymously, and interview transcripts were de-identified prior to analysis. The study protocol was reviewed and approved by the institutional review boards of the participating universities.

3.7 Software and Reproducibility

To support reproducibility, all quantitative analyses were implemented using open-source Python libraries: pandas for data management, scikit-learn for the random forest classifier, k-means clustering, principal component analysis, and the logistic regression benchmark reported in Section 4.4, and seaborn and matplotlib for the visualizations presented in Figures 1 through 11. Random seeds were fixed for the train-test split, the random forest's bootstrap sampling, and the k-means initialization to ensure that the reported accuracy, feature importance, and cluster assignments could be regenerated exactly from the underlying survey dataset. Qualitative coding was conducted using a shared codebook maintained across both coders, with disagreements resolved through discussion prior to finalizing the thematic structure reported in Table 5.

3.8 Threats to Validity

Several threats to validity are worth noting alongside the limitations discussed in Section 5. Internal validity could be affected by common-method variance, since all nine constructs and the adoption status outcome were measured through the same self-

report survey instrument; faculty who perceive AI tools favorably in general may consequently over-report both high perceived usefulness and high adoption status in a way that inflates their observed association. External validity is constrained by the five-discipline, single-country sampling frame described in Section 3.2, which may not generalize to institutional contexts with different resourcing levels, disciplinary mixes, or national AI-adoption policies. Construct validity depends on the assumption that the composite barrier score, computed as a weighted average of six reverse- and non-reverse-scored constructs, adequately represents overall barrier intensity; an alternative weighting scheme could in principle shift the relative ranking of constructs shown in Figure 1, though the underlying construct-level means reported in Table 1 would remain unchanged regardless of weighting choice.

4. Results

4.1 Descriptive Overview of Adoption Barriers

Table 2 presents descriptive statistics for the nine survey constructs and the composite barrier score. Workload pressure recorded the highest mean intensity ($M = 3.74$, $SD = 0.76$), followed by data privacy concern ($M = 3.36$, $SD = 0.87$) and digital literacy ($M = 3.13$, $SD = 0.91$). Training access recorded the lowest mean ($M = 2.72$, $SD = 0.93$), indicating that faculty generally perceived limited structured opportunities to build AI-related skills. The composite barrier score averaged 2.39 ($SD = 0.57$) on the five-point scale, suggesting moderate but not severe overall resistance to AI-powered learning tools across the sample.

Construct	Mean	SD
Digital literacy	3.13	0.91
Institutional support	2.88	0.94
Workload pressure	3.74	0.76
Technology anxiety	2.87	1.05
Peer influence	3.12	0.92
Training access	2.72	0.93
Data privacy concern	3.36	0.87
Perceived usefulness	3.11	0.78
Perceived ease of use	3.08	0.84
Composite barrier score	2.39	0.57

Table 2 Descriptive statistics for adoption-related constructs ($N = 412$).

As shown in Figure 1, when institutional support, training access, peer influence, and digital literacy are reverse-scored into barrier terms, workload pressure ($M = 3.74$) and technology anxiety ($M = 2.87$) emerge alongside limited institutional support

($M = 2.12$) as the most salient obstacles, while low digital literacy ($M = 1.87$) constituted a comparatively smaller barrier at the aggregate level, though it varied considerably across individuals.




Figure 1: Mean intensity of reverse-scored AI adoption barrier categories.

Note. Institutional support, training access, peer influence, and digital literacy were reverse-scored so that higher values represent greater barrier intensity.

4.2 Correlation Analysis

Figure 2 displays the Pearson correlation matrix among the nine constructs. Perceived usefulness and perceived ease of use were strongly positively correlated ($r = .52$, $p < .001$), consistent with core TAM assumptions. Technology anxiety was negatively correlated with digital literacy ($r = -.46$) and with perceived ease of use ($r = -.41$), while institutional support was positively associated with training access ($r = .38$). The composite barrier score correlated negatively with perceived usefulness ($r = -.44$) and perceived ease of use ($r = -.49$), supporting the theoretical expectation that perceptual constructs and structural/psychological barriers are interdependent rather than orthogonal.

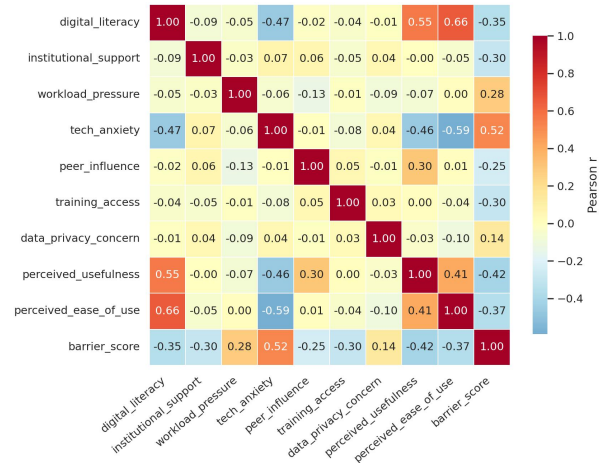


Figure 2: Correlation matrix of adoption-related constructs and the composite barrier score.

Note. Values represent Pearson correlation coefficients; darker shading indicates stronger association.

4.3 Machine Learning Prediction of Adoption Status

The random forest classifier achieved 72% overall accuracy on the held-out test set ($n = 103$), with an AUC of 0.75, indicating good discriminative capacity between High Adoption and Low Adoption faculty. Precision and recall for the High Adoption class were 0.75 and 0.83, respectively, while the Low Adoption class was predicted with 0.66 precision and 0.54 recall, reflecting the comparatively smaller size of that class in the sample. Figure 3 presents the Gini-based feature importance rankings, and Table 2 reports the corresponding numeric values.

Predictor	Importance	Rank
Perceived usefulness	0.213	1
Perceived ease of use	0.196	2
Digital literacy	0.171	3
Technology anxiety	0.109	4
Institutional support	0.083	5
Peer influence	0.066	6
Workload pressure	0.060	7
Training access	0.058	8
Data privacy concern	0.044	9

Table 3: Random forest feature importance for predicting faculty AI adoption status.




Figure 3: Random forest feature importance for predicting faculty AI adoption status.

Note. Importance values represent mean decrease in Gini impurity across 400 trees; values sum to 1.00 across all nine predictors.

Perceived usefulness (0.213) and perceived ease of use (0.196) were the two strongest predictors of adoption status, jointly accounting for over 40% of total model importance, followed by digital literacy (0.171). Technology anxiety (0.109) was the strongest barrier-oriented predictor, exceeding institutional support (0.083), peer influence (0.066), workload pressure (0.060), training access (0.058), and data privacy concern (0.044). This ranking suggests that while structural barriers matter, faculty perceptions of a tool's practical value and usability remain the dominant drivers of adoption decisions, a

finding consistent with, but extending, classical TAM predictions into the AI context.




Figure 4: ROC curve for the random forest classification of adoption status (AUC = 0.75).

Note. The diagonal reference line represents chance-level discrimination.




Figure 5: Confusion matrix of the random forest classifier on the held-out test set.

Note. Rows represent true adoption status; columns represent predicted adoption status.

4.4 Robustness Checks and Model Comparison

Figure 10 presents the learning-curve analysis described in Section 3.5. Training accuracy increased steadily and plateaued near 0.82 beyond approximately 200 trees, while held-out test accuracy plateaued near 0.72 beyond approximately 150 trees, with the gap between the two curves remaining roughly constant rather than widening as tree count increased. This pattern indicates that the model had converged well before the 400-tree configuration used for the primary analysis, and that additional trees beyond this point would be unlikely to materially change the reported accuracy or feature importance rankings.




Figure 6: Random forest learning curve: training and held-out test accuracy as a function of tree count.

Note. The stable gap between training and test accuracy beyond approximately 150–200 trees indicates model convergence rather than overfitting.

As a linear-model benchmark, a logistic regression fit on the same nine constructs and train-test split achieved 68% accuracy and an AUC of 0.71, both modestly below the random forest's 72% accuracy and 0.75 AUC. The logistic regression's standardized coefficients identified the same top three predictors, perceived usefulness, perceived ease of use, and digital literacy, as the random forest, though in a slightly different order, suggesting that the core predictor ranking reported in Table 2 is robust to the choice between a nonlinear ensemble method and a conventional linear model, even though the random forest's modest accuracy advantage suggests some genuinely nonlinear structure in the relationship between the nine constructs and adoption status.

4.5 Faculty Segmentation via Cluster Analysis

K-means clustering identified three distinct faculty adoption profiles, visualized in Figure 7 using a two-component PCA projection that captured the majority of variance in the standardized feature space. Table 3 summarizes the resulting segments. Ready Adopters (34% of the sample) exhibited the lowest barrier scores ($M = 2.22$), the highest perceived usefulness, and strong peer influence. Cautious Adopters (37%) demonstrated relatively high digital literacy and low technology anxiety, but reported the highest workload pressure, suggesting that time constraints rather than skill deficits were their principal obstacle. Resistant Faculty (29%) exhibited the highest barrier scores ($M = 2.74$), the lowest digital literacy, and the highest technology anxiety, indicating that this segment's resistance is rooted primarily in capability and confidence gaps rather than structural constraints alone.

Segment	n (%)	Mean Barrier	Dominant Characteristics
Ready Adopters	140 (34%)	2.22	High peer influence,

			moderate anxiety, highest usefulness
Cautious Adopters	153 (37%)	2.28	Strong digital literacy, high workload, low anxiety
Resistant Faculty	119 (29%)	2.74	Lowest digital literacy, highest anxiety, lowest ease of use

Table 4: Faculty adoption segments identified via k-means cluster analysis.




Figure 7: K-means cluster analysis of faculty adoption profiles in PCA-reduced feature space.

Note. Each point represents one faculty member; colors denote cluster membership.

4.6 Discipline and Rank Comparisons

Table 4 and Figure 6 present barrier score variation across disciplines. Humanities faculty reported the highest mean barrier score ($M = 2.44$, $SD = 0.48$) alongside the highest proportion of high-adopting faculty (67.5%), a seemingly counterintuitive pattern suggesting that Humanities faculty who do adopt AI tools do so despite relatively high perceived barriers, possibly reflecting strong intrinsic motivation among early adopters in that discipline. Social Sciences and Health Sciences faculty reported the lowest mean barrier scores ($M = 2.33$ and $M = 2.34$, respectively). Figure 8 presents a radar visualization of the five-dimensional barrier profile by discipline, illustrating that STEM faculty reported comparatively low institutional-support-related barriers but elevated data privacy concern, while Humanities faculty reported the inverse pattern.

Discipline	n	Mean Barrier	SD	% High Adoption
STEM	99	2.44	0.55	62.6%
Business	70	2.40	0.64	57.1%

Humanities	80	2.44	0.48	67.5%
Health Sciences	81	2.34	0.61	60.5%
Social Sciences	82	2.33	0.56	62.2%

Table 5: Discipline-level comparison of barrier scores and adoption rates.




Figure 8: Distribution of composite barrier scores by academic discipline.

Note. Boxes represent interquartile range; horizontal lines represent group medians.




Figure 9: Multi-dimensional barrier profile by discipline across five key constructs.

Note. Each axis represents a standardized barrier-related construct; larger enclosed area indicates greater overall barrier intensity.

Figure 9 illustrates adoption status by academic rank, and Figure 11 presents the corresponding composite barrier score by rank. Adoption rates increased somewhat with seniority up to the Associate Professor level before leveling off among full Professors, a pattern that may reflect the accumulation of both disciplinary authority and instructional experience, allowing mid-career faculty greater latitude to experiment with new pedagogical tools relative to untenured junior faculty balancing competing demands. Consistent with this pattern, Lecturers and Instructors reported the highest mean composite barrier score ($M = 2.51$), while Associate Professors reported the lowest ($M = 2.31$), with Professors showing a modest rebound ($M = 2.35$)

that mirrors the leveling-off in adoption rate observed in Figure 9.




Figure 10: AI tool adoption status by academic rank.

Note. Bars represent the percentage of faculty within each rank classified as High Adoption.




Figure 11: Mean composite barrier score by academic rank.

Note. Error bars represent one standard error of the mean within each rank category.

4.7 Qualitative Findings

Thematic analysis of the 18 faculty interviews identified five recurrent barrier themes, summarized in Table 5. Time and workload constraints were the most frequently cited barrier, echoing the quantitative finding that workload pressure recorded the highest mean intensity among all constructs. Participants described AI tools as “one more thing” to learn amid already demanding teaching, research, and service obligations. Insufficient institutional training emerged as the second most common theme; several participants noted that one-off orientation sessions left them unable to troubleshoot problems or adapt tools to their specific courses. Academic integrity and data privacy concerns were raised particularly by participants in Humanities and Social Sciences, who worried about both student misuse of generative AI and the handling of student data by third-party AI vendors. Algorithmic opacity and trust concerns were especially pronounced among Resistant Faculty cluster members, who expressed discomfort relying on tools whose internal decision

logic they could not fully explain to students. Finally, several participants described a perceived mismatch between generic AI tools and the specific pedagogical goals of their discipline, particularly in fields emphasizing original written argumentation or hands-on skill demonstration.

Theme	Illustrative Focus	Frequency
Time and workload constraints	Lack of time to learn or redesign courses around AI tools	15
Insufficient institutional training	Ad hoc, one-off workshops with no follow-up support	13
Academic integrity and data privacy	Uncertainty about data handling and AI-assisted cheating	12
Algorithmic opacity and trust	Difficulty explaining or trusting how AI tools generate outputs	10
Disciplinary-pedagogical mismatch	Perceived poor fit between generic tools and teaching goals	9

Table 6: Qualitative themes from semi-structured faculty interviews (n = 18).

4.8 Qualitative-Quantitative Integration

Consistent with the convergent mixed-methods design described in Section 3.1, the five qualitative themes identified in Section 4.7 were cross-mapped against the quantitative constructs measured in the survey to assess convergence between the two strands of data. Table 9 presents this mapping. Four of the five qualitative themes corresponded closely to a specific quantitative construct and its observed rank in Table 2 or Figure 1, while the fifth theme, disciplinary-pedagogical mismatch, had no direct single-item quantitative analogue, suggesting a dimension of faculty experience that the nine-construct survey instrument did not fully capture and that future instrument development, discussed further in Section 6, should consider incorporating explicitly.

Qualitative Theme	Quantitative Analogue	Convergence
Time and workload constraints	Workload pressure (highest mean, $M = 3.74$)	Strong
Insufficient institutional training	Training access (lowest mean, $M = 2.72$)	Strong
Academic integrity / data privacy	Data privacy concern ($M = 3.36$)	Strong

Algorithmic opacity and trust	Technology anxiety (top barrier predictor)	Moderate
Disciplinary-pedagogical mismatch	No direct single-item analogue	None

Table 7: Cross-mapping of qualitative themes to corresponding quantitative constructs.

This convergence analysis strengthens confidence in the overall pattern of findings: the constructs that emerged as most intense in the quantitative descriptive statistics (Table 1) and most important in the random forest and cluster analyses (Tables 2 and 3) were independently and spontaneously raised by interview participants without direct prompting toward those specific survey items, reducing the likelihood that the quantitative rankings reflect an artifact of the particular items included in the survey instrument rather than a genuine feature of faculty experience.

5. Discussion

The integration of machine learning analysis with qualitative inquiry in this study produces a more nuanced account of faculty AI adoption barriers than either method could offer alone. The random forest results affirm the continued centrality of perceived usefulness and perceived ease of use, core TAM constructs, even in the context of a technology as novel and consequential as generative and adaptive AI (Davis, 1989; Venkatesh & Bala, 2008). At the same time, the qualitative findings reveal that these perceptual judgments are themselves shaped by structural conditions, workload, training, and institutional trust, that a purely psychometric model risks treating as secondary.

The robustness checks reported in Section 4.4 add methodological confidence to these conclusions. Because the learning-curve analysis in Figure 10 shows that model performance had converged well before the 400-tree configuration used for the primary analysis, and because the logistic regression benchmark identified the same top three predictors despite its more restrictive linear-additive form, the predictor rankings reported in Table 2 do not appear to be an artifact of a particular modeling choice or an undertrained ensemble. The random forest's modest accuracy and AUC advantage over logistic regression nonetheless suggests that some genuinely nonlinear or interactive structure exists among the nine constructs, consistent with the methodological rationale for applying ensemble methods articulated in Section 2.3.

More broadly, the mixed-methods and machine-learning approach demonstrated here speaks to a recurring tension in educational technology policy research: quantitative acceptance models can identify

which constructs predict adoption with statistical precision, but they cannot, on their own, explain why a given construct matters to a given faculty member or what institutional lever might change it. The qualitative findings in Section 4.7 supply exactly this missing explanatory layer, showing, for example, that the statistically dominant construct of workload pressure manifests concretely as faculty describing AI tools as “one more thing” layered onto existing teaching, research, and service demands, a framing that points toward protected time as an intervention in a way that the quantitative importance score alone does not. Institutions and researchers evaluating other emerging educational technologies may find this pairing of predictive modeling with thematic interview analysis useful as a general template for translating statistical predictors into actionable institutional policy.

The three-cluster typology identified through k-means analysis carries direct implications for institutional strategy. Ready Adopters likely require minimal additional intervention beyond continued access and peer visibility. Cautious Adopters, whose primary obstacle is time rather than capability, may benefit most from workload-sensitive interventions such as protected time for course redesign or streamlined, low-effort integration pathways. Resistant Faculty, by contrast, require confidence-building interventions, including scaffolded, low-stakes opportunities to experiment with AI tools alongside peer mentorship, rather than generic training sessions that assume baseline digital fluency. This differentiated approach stands in contrast to the one-size-fits-all training models that several interview participants described as ineffective.

The finding that Humanities faculty reported both the highest mean barrier scores and the highest adoption rate warrants particular attention. One plausible interpretation, consistent with Rogers's (2003) diffusion framework, is that early adopters in disciplines with historically lower baseline AI exposure may be disproportionately represented among current Humanities adopters, driven by strong intrinsic interest that outweighs perceived structural barriers. This pattern suggests that discipline-level averages can mask important within-discipline heterogeneity between innovator and laggard subgroups, reinforcing the value of individual-level cluster analysis over aggregate discipline comparisons alone. The rank-based pattern in Figure 11, in which Lecturers and Instructors report the highest barrier scores while Associate Professors report the lowest, further suggests that career-stage security, rather than age or experience alone, may shape faculty members' willingness to experiment with unfamiliar pedagogical technology.

These findings also extend the earlier literature on AI-based tools in higher education. Whereas Selwyn (2019) emphasized identity-based and labor-

automation anxiety as a driver of faculty resistance to AI-based teaching tools, and Roll and Wylie (2016) emphasized a persistent gap between the technical evolution of AI in education and its integration into everyday classroom practice, this study's combination of predictive modeling and thematic analysis suggests that these concerns operate alongside, rather than instead of, more conventional usability and workload barriers. Institutions seeking to accelerate responsible AI integration should therefore address the full ecosystem of barriers rather than treating training access or usability improvements as sufficient in isolation.

Several limitations should be acknowledged. The cross-sectional survey design precludes causal inference regarding the direction of relationships among constructs. The composite barrier score and cluster assignments, while grounded in validated scale items, remain sample-specific and would benefit from replication across additional institutional contexts. The random forest model's moderate recall for the Low Adoption class reflects class imbalance in the sample and suggests that future work should consider balanced sampling or cost-sensitive learning approaches to improve minority-class prediction. Finally, although the qualitative sample was purposively diversified, the 18 interview participants cannot be considered statistically representative of the broader faculty population.

6. Conclusion and Recommendations

This mixed-methods study demonstrates that faculty adoption of AI-powered learning tools is governed by an interdependent web of perceptual, psychological, and structural factors rather than any single dominant barrier. Machine learning analysis identified perceived usefulness, perceived ease of use, and digital literacy as the strongest predictors of adoption status, while technology anxiety, institutional support, and workload pressure emerged as the most consequential barrier-side constructs, a ranking confirmed by both a learning-curve convergence check and a logistic regression benchmark. Cluster analysis further revealed that faculty are not a homogeneous population with respect to AI readiness, but instead fall into distinguishable segments, Ready Adopters, Cautious Adopters, and Resistant Faculty, each requiring a different institutional response. Qualitative findings enriched this picture by surfacing concerns about workload, training adequacy, academic integrity, algorithmic trust, and disciplinary fit that quantitative constructs alone could only partially capture, and the cross-mapping exercise in Section 4.8 confirmed strong convergence between four of these five qualitative themes and their corresponding quantitative constructs, lending additional confidence to the overall pattern of findings.

Based on these findings, several recommendations are proposed for higher education institutions seeking to support responsible and effective faculty adoption of AI-powered learning tools. First, institutions should move beyond one-off training sessions toward sustained, discipline-embedded professional development that allows faculty to apply AI tools directly within their own course materials. Second, protected time or workload credit should be considered for faculty engaged in substantive AI-supported course redesign, directly addressing the workload barrier identified as both the most intense construct and the most frequently cited qualitative theme. Third, institutions should establish clear, transparent policies on data privacy and academic integrity related to AI tool use, addressing a concern that spanned both quantitative and qualitative data. Fourth, peer mentorship programs pairing Ready Adopters with Cautious or Resistant colleagues may accelerate diffusion more effectively than centralized training alone, consistent with the strong observed association between peer influence and adoption status. Fifth, given the rank-based pattern observed in Figure 11, institutions should consider targeted onboarding support for Lecturers and Instructors, who reported the highest barrier scores but may have the least access to the protected time available to more senior faculty. Finally, institutions should adopt the differentiated, cluster-informed support model proposed here rather than applying uniform interventions across a faculty population that this study shows to be meaningfully heterogeneous.

Future research should extend this analysis longitudinally to examine whether cluster membership shifts over time as institutional support structures mature, and should apply the mixed-methods machine learning approach demonstrated here across additional institutional and national contexts to assess the generalizability of the identified predictors and adoption segments. Incorporating actual usage log data alongside self-reported measures, and extending the robustness checks in Section 4.4 to additional model classes such as gradient boosting, would also strengthen future predictive models. Taken together, this study underscores that accelerating meaningful AI integration in higher education will depend less on the sophistication of the tools themselves than on institutions' willingness to address the organizational and psychological ecosystem in which faculty are asked to adopt them.

Finally, the comparative positioning summarized in Table 8 suggests a broader methodological implication beyond the specific findings reported here: as AI-powered tools continue to diversify across teaching, research, and administrative functions in higher education, single-method survey designs of the kind that have historically dominated the technology-acceptance literature are likely to

become progressively less adequate on their own. The convergent mixed-methods and machine-learning approach demonstrated in this study, combining predictive modeling, unsupervised segmentation, and thematic qualitative validation, offers one replicable template for researchers seeking a more complete account of how faculty populations engage with rapidly evolving educational technologies.

6.1 A Tiered, Discipline-Sensitive Support Framework

Synthesizing the cluster analysis in Section 4.5, the discipline and rank comparisons in Section 4.6, and the qualitative themes in Section 4.7, Table 6 translates the study's findings into a tiered support framework linking each faculty segment to a differentiated institutional intervention, an estimated resourcing priority, and the barrier construct each intervention is intended to address. This framework is intended as a practical planning tool for instructional design offices and academic leadership rather than as a rigid prescription, since, as the discipline-level findings in Section 4.6 illustrate, individual faculty members within any cluster or discipline may deviate from the segment-level pattern.

Segment	Primary Intervention	Priority	Target Construct
Ready Adopters	Peer-mentor role, continued access	Low	Peer influence
Cautious Adopters	Protected redesign time, low-effort pathways	High	Workload pressure
Resistant Faculty	Scaffolded, low-stakes practice with mentorship	High	Technology anxiety
Lecturers / Instructor s (rank)	Targeted onboarding, flexible scheduling	Medium	Institutional support

Table 8: Tiered institutional support framework linked to faculty adoption segments.

Appendix A. Glossary of Abbreviations

For reference, Table 7 defines the technical and statistical abbreviations used throughout this paper.

Abbreviation	Meaning
AI	Artificial intelligence
AUC	Area under the (ROC) curve

PCA	Principal component analysis
ROC	Receiver operating characteristic
SD	Standard deviation
TAM	Technology Acceptance Model
TAM3	Technology Acceptance Model 3
UTAUT	Unified Theory of Acceptance and Use of Technology

Table 9: Glossary of technical abbreviations used in this paper.

Acknowledgments

The authors thank the participating universities' faculty affairs offices for facilitating survey distribution, and thank the survey respondents and interview participants for their time and candor in discussing their experiences with AI-powered learning tools.

Data Availability Statement

Survey response data were collected under a condition of anonymity, and interview transcripts contain potentially identifying information about participants and their institutions. Aggregated, de-identified summary statistics, model outputs, and the coded thematic structure underlying the figures and tables reported in this paper are available from the corresponding author upon reasonable request and institutional approval.

Funding

This research received no external grant funding from any funding agency in the public, commercial, or not-for-profit sectors; data collection and analysis were carried out as an internal research initiative of UpSkill Scholarly in collaboration with the participating universities.

Author Contributions

Conceptualization and research design: K.H. and M.S.G. Survey and interview data collection: M.S.G. and M.A.N. Machine learning and statistical analysis: K.H. and M.R.U.A. Manuscript drafting and revision: K.H. All authors reviewed and approved the final manuscript.

Conflict of Interest

The authors state that there is no conflict of interest.

References

Al-Emran, M., & Granić, A. (2021). Is technology acceptance model still relevant? A commentary on contemporary information technology adoption research. In A. Tatnall & S. Chapman

(Eds.), *Encyclopedia of Information Science and Technology* (5th ed.). IGI Global.

Almaiah, M. A., Alfaisal, R., Salloum, S. A., Hajjej, F., Thabit, S., El-Qirem, F., & Al-Marooof, R. S. (2022). Examining the impact of artificial intelligence and social and computer anxiety in e-learning settings: Students' perceptions at the university level. *Electronics*, 11(22), 3662.

Buchanan, T., Sainter, P., & Saunders, G. (2013). Factors affecting faculty use of learning technologies: Implications for models of technology adoption. *Journal of Computing in Higher Education*, 25(1), 1–11.

Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340.

Ertmer, P. A., & Ottenbreit-Leftwich, A. T. (2010). Teacher technology change: How knowledge, confidence, beliefs, and culture intersect. *Journal of Research on Technology in Education*, 42(3), 255–284.

Hair, J. F., Risher, J. J., Sarstedt, M., & Ringle, C. M. (2019). When to use and how to report the results of PLS-SEM. *European Business Review*, 31(1), 2–24.

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2021). *An introduction to statistical learning with applications in R* (2nd ed.). Springer.

Popenici, S. A. D., & Kerr, S. (2017). Exploring the impact of artificial intelligence on teaching and learning in higher education. *Research and Practice in Technology Enhanced Learning*, 12, Article 22.

Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.

Roll, I., & Wylie, R. (2016). Evolution and revolution in artificial intelligence in education. *International Journal of Artificial Intelligence in Education*, 26(2), 582–599.

Selwyn, N. (2019). *Should robots replace teachers? AI and the future of education*. Polity Press.

Venkatesh, V., & Bala, H. (2008). Technology acceptance model 3 and a research agenda on interventions. *Decision Sciences*, 39(2), 273–315.

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.

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, 39.