Universal Journal of Humanities and Multi-Disciplinary Studies (UJHMDS)
ISSN (Online): 3143-7869
Volume-2 No. 3 || Special Issue Jul – Sep 2026 || PP. 95 – 106
https://ujhmds.org/

Artificial Intelligence Tools and Personalized Instruction in Higher Education

Dr. Pushpa Sharma

Associate Professor

Mansa College of Education

Kohka, Kurud, Chhattisgarh

http://www.ujhmds.org/ || Vol. 2 No. 3 (2026): Special August Issue

Date of Submission: 02-07-2026	**Date of Acceptance:** 22-07-2026	**Date of Publication:** 17-08-2026

Abstract— Artificial intelligence is reshaping personalized instruction in higher education by enabling adaptive content delivery, automated feedback, learner profiling, and conversational academic support. However, existing research frequently examines adaptive learning platforms, predictive analytics, or generative AI tools as separate interventions rather than as components of an integrated instructional system. This study addresses that gap by proposing a human-in-the-loop personalization framework in which machine-learning-based learner profiling, generative AI-supported instructional adaptation, and instructor validation operate as interconnected processes. The central objective is to investigate whether coordinated AI personalization can improve learning achievement, engagement, instructional relevance, and learner autonomy without reducing meaningful instructor involvement. The proposed research conceptualizes personalization as a dynamic process informed by prior achievement, interaction behaviour, learning pace, feedback responses, and evolving competence rather than as static content recommendation.

Keywords— Artificial Intelligence, Personalized Instruction, Higher Education, Generative AI, Learning Analytics, Adaptive Learning

Introduction

Higher education has entered a period in which artificial intelligence is moving from specialized educational technology laboratories into routine teaching and learning environments. Machine learning, learning analytics, intelligent tutoring systems, recommender systems, natural language processing, and more recently generative artificial intelligence have substantially expanded the technical possibilities for adapting educational experiences to individual learners. The rapid growth of large language model-based applications since 2022 has further altered this landscape because students can now interact directly with systems capable of generating explanations, examples, questions, summaries, feedback, and alternative representations of academic content. Consequently, personalized instruction is no longer confined to rule-based adaptive platforms in which students are automatically directed through predetermined learning paths.

The contemporary concept of AI-supported personalization involves a considerably wider instructional architecture. Student activity within a learning management system can be analysed to estimate competence, identify recurring errors, detect patterns of disengagement, and recommend resources. Predictive models can estimate whether students require additional academic support, while generative AI can reformulate explanations according to learner proficiency or produce differentiated formative activities. Conversational AI can additionally provide immediate academic assistance outside scheduled instructional hours. These capabilities create the possibility of a learning environment that adjusts

@2026. This is an open access article distributed under the terms of the Creative Commons License [CC BY NC 4.0] and is available on www.ujhmds.org

Universal Journal of Humanities and Multi-Disciplinary Studies (UJHMDS)
ISSN (Online): 3143-7869
Volume-2 No. 3 || Special Issue Jul – Sep 2026 || PP. 95 – 106
https://ujhmds.org/

not merely to what students have completed, but to how they are progressing and where instructional intervention is likely to be most valuable.

The OECD has identified personalization as one of the important potential benefits of digital transformation in education and has emphasized that recent generative AI developments could enable learning support adapted to individual needs. At the same time, the organization stresses that teachers remain central to educational transformation and that data governance, privacy, interoperability, and professional capability are necessary for responsible implementation. These observations are particularly relevant to universities, where personalization must accommodate considerable differences in prior knowledge, language proficiency, digital competence, academic motivation, and self-regulated learning ability.

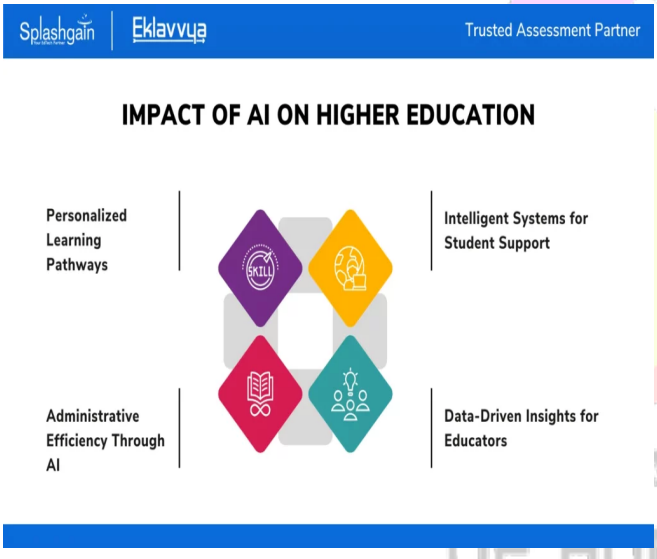


Fig. 1: The Impact of AI on Higher Education Learning and Teaching

Source: *https://www.eklavvya.com/blog/ai-higher-education-trend/*

Research on artificial intelligence in higher education has expanded considerably, although the evidence base remains uneven. Bond et al. (2024), in a meta-systematic review of AI in higher education, called for stronger methodological rigor, ethical analysis, and interdisciplinary collaboration. Their synthesis also illustrates the diversity of applications grouped under artificial intelligence, including prediction, intelligent tutoring, adaptive systems, assessment, and personalization. Zawacki-Richter et al. (2024) similarly observed that considerable uncertainty remains regarding the actual influence of contemporary AI on higher-education institutions, teaching processes, students, and instructors. They further noted that educational research has not always developed at the same rate as advances in AI algorithms and applications.

The introduction of generative AI has intensified this research problem. Conventional adaptive systems generally select among existing resources or modify difficulty levels according to performance indicators. Generative systems can instead create new instructional material dynamically. A student struggling with a statistical concept, for example, could receive a simplified explanation, a discipline-specific analogy, a worked example, and progressively more difficult practice questions generated for the immediate learning context. Another student demonstrating mastery of the same concept could receive an advanced application task rather than repeating elementary material. Such responsiveness offers a fundamentally different form of instructional differentiation.

Yet technical adaptability is not equivalent to pedagogically valid personalization. AI-generated explanations may contain inaccuracies, present unsuitable cognitive difficulty, provide excessive assistance, or inadvertently encourage dependence on generated solutions. Research on early educational use of ChatGPT has consequently identified both possibilities and risks. Tlili et al. (2023) documented extensive educational interest in conversational AI while emphasizing the importance of human–machine collaboration and careful consideration of ethical and pedagogical implications. Related empirical work examining early adopters reported concerns that excessive reliance on generative AI could encourage superficial learning and weaken critical or social dimensions of learning.

Research Gap

A significant limitation in the present literature is that the components of AI personalization are commonly studied independently. Predictive analytics research frequently concentrates on identifying at-risk students; intelligent tutoring studies evaluate automated guidance; recommender-system studies examine content selection; and recent generative AI research often investigates acceptance, perceptions, academic integrity, or generic assistance. Recent reviews confirm that AI-supported personalized learning remains heterogeneous in its terminology, technologies, pedagogical foundations, and empirical evaluation. A 2025 systematic review of AI-supported personalized learning in higher education found that quantitative approaches and supervised machine-learning methods remain prominent while explicit pedagogical grounding and systematic consideration of equity are comparatively limited. A 2026 mapping review likewise noted conceptual fragmentation between individualized teaching, personalized learning, adaptive learning, and AI-supported instructional environments.

@2026. This is an open access article distributed under the terms of the Creative Commons License [CC BY NC 4.0] and is available on www.ujhmds.org

Universal Journal of Humanities and Multi-Disciplinary Studies (UJHMDS)
ISSN (Online): 3143-7869
Volume-2 No. 3 || Special Issue Jul – Sep 2026 || PP. 95 – 106
https://ujhmds.org/

Less attention has therefore been given to a complete instructional cycle in which **learner modelling predicts instructional need, generative AI adapts learning materials, student interaction produces new behavioural evidence, and instructors validate or override consequential AI recommendations**. Moreover, many studies determine effectiveness primarily through examination scores, acceptance measures, or satisfaction ratings. These indicators cannot establish whether the personalization mechanism itself made appropriate decisions.

The present study therefore introduces a second evaluative layer termed **Personalization Alignment Quality (PAQ)**. Rather than asking only whether AI-supported students perform better, PAQ evaluates whether the instructional adaptation corresponded appropriately to diagnosed learner needs. For example, the system may recommend remedial material to a student whose subsequent performance indicates that enrichment would have been more appropriate. A conventional achievement comparison would not identify this personalization error. The proposed approach consequently treats instructional effectiveness and personalization correctness as related but distinct outcomes.

A further gap concerns instructor agency. Considerable debate frames educational AI through a binary question of whether AI should or should not perform particular teaching functions. The more practically relevant question is how responsibility should be distributed between AI and academic staff. A human-in-the-loop design provides a mechanism for examining this issue empirically. Instructor overrides can themselves become data: frequent rejection of particular recommendation types may indicate systematic weaknesses in the learner model or generative adaptation mechanism.

Accordingly, the main objective of the proposed study is to evaluate whether a human-supervised AI personalization architecture can generate measurable improvements in academic achievement, engagement, and instructional relevance while maintaining learner autonomy and instructor oversight. The study will examine not only whether the system produces better outcomes but also **when algorithmic recommendations are appropriate, when instructors intervene, and whether those interventions improve subsequent learning trajectories**.

Literature Review

Artificial Intelligence and the Evolution of Personalized Learning

Personalized learning predates contemporary artificial intelligence. Earlier educational technologies used branching rules, mastery thresholds, computer-assisted instruction, and intelligent tutoring systems to modify instructional sequences. The expansion of educational data infrastructures subsequently enabled more sophisticated learner modelling. Machine-learning approaches can analyse multidimensional information such as assessment history, resource access frequency, submission behaviour, response time, discussion participation, and navigation sequences to infer patterns not readily detectable through manual observation.

Recent developments extend this functionality from retrospective analysis to adaptive intervention. Artificial intelligence can potentially connect diagnosis and instructional action: a model identifies a weakness, a recommendation engine selects an intervention, and a generative component modifies the instructional presentation. The crucial research issue is whether this technical pipeline corresponds to legitimate pedagogical decisions.

Systematic evidence indicates substantial interest in AI-enabled personalized learning but also significant conceptual diversity. Peng and Li's recent systematic review identified personalization, adaptive feedback, learner support, intelligent tools, and algorithmic techniques across higher-education research while observing that pedagogical theories are not always explicitly integrated into system design. This is consequential because an algorithm optimized exclusively for short-term prediction accuracy could recommend activities that improve immediate test performance while doing little to develop durable understanding, self-regulation, or transfer.

The present study responds to this issue by positioning the personalization engine as an instructional decision-support system rather than an autonomous optimizer. The underlying assumption is that student-data patterns provide useful evidence but not complete explanations of learning. An extended period spent viewing a resource, for example, may reflect difficulty, deep engagement, distraction, or an interruption unrelated to learning. Consequently, algorithmic indicators require contextual interpretation.

Generative AI as an Instructional Adaptation Mechanism

Large language models have altered the technological feasibility of personalization because instructional variation can now be generated in real time. Previously, adaptive platforms depended heavily on repositories containing pre-authored explanations, questions, and remediation pathways. Generative AI reduces this constraint by producing alternative examples, explanations, practice exercises, hints, or summaries on demand.

Systematic reviews of generative AI for teaching and learning demonstrate a growing range of educational applications, but they also identify issues concerning reliability, ethics, academic integrity, and pedagogically meaningful use. The distinction between content generation and learning support is therefore critical. Generating more content does not necessarily create better instruction. Effective personalization

@2026. This is an open access article distributed under the terms of the Creative Commons License [CC BY NC 4.0] and is available on www.ujhmds.org

Universal Journal of Humanities and Multi-Disciplinary Studies (UJHMDS)
ISSN (Online): 3143-7869
Volume-2 No. 3 || Special Issue Jul – Sep 2026 || PP. 95 – 106
https://ujhmds.org/

requires information about the learner, an explicit instructional objective, appropriate difficulty calibration, and safeguards against erroneous or excessive assistance.

Evidence from university-student use of generative AI further suggests variation in how learners interact with such systems. A recent systematic literature review found differences in usage patterns and raised the possibility that AI could amplify disparities when learners possess different levels of competence in critically using generated outputs. Personalized systems must therefore account for AI literacy as well as subject knowledge. A student who accepts generated content uncritically requires a different instructional design from a student capable of evaluating, revising, and challenging AI responses.

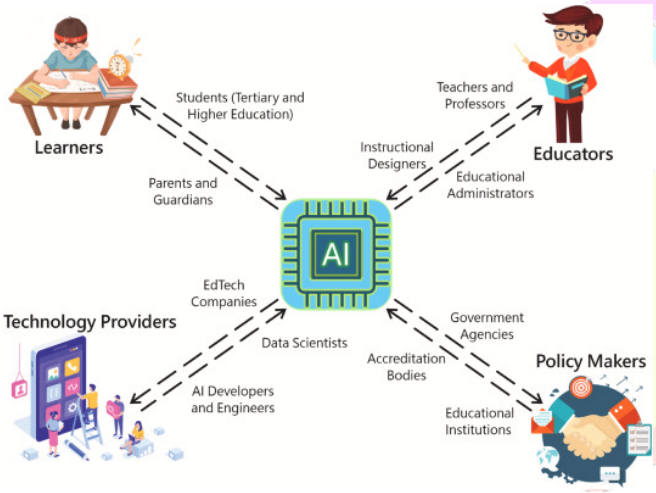


Fig. 2: Artificial Intelligence in Personalized Learning

Source: *https://www.sciencedirect.com/science/article/pii/S2590291125008447*

Learning Analytics and Dynamic Learner Modelling

Learning analytics provides the data foundation for computational personalization. Traditional personalization may rely on a limited set of relatively static variables, such as prior grades or a diagnostic test. Dynamic personalization instead requires continuous updating as students interact with course material.

A learner model can incorporate assessment performance, error patterns, time-on-task indicators, activity completion sequences, response latency, formative feedback usage, and interaction frequency. These variables can support estimation of mastery or the probability that particular forms of intervention will be beneficial. Nevertheless, educational prediction should not be confused with causal explanation. A model may accurately predict poor performance without identifying why a student is struggling.

This distinction explains why the present research separates **learner-state estimation** from **instructional adaptation**. The first process identifies a probable learning condition, while the second determines an educational response. A high-risk prediction does not itself indicate whether the student requires conceptual remediation, additional examples, motivational support, reduced complexity, or instructor consultation.

The proposed architecture consequently treats personalization as a sequential decision process. Student behaviour updates the learner profile; the system produces an instructional recommendation; the recommendation is implemented or modified; and subsequent behaviour provides evidence about whether the intervention was successful. This creates a closed feedback loop that can be evaluated experimentally.

Human Oversight and Instructor Agency

The role of instructors represents one of the most important yet insufficiently operationalized dimensions of educational AI. OECD analysis emphasizes that teachers must remain central actors in digital transformation and that effective AI adoption requires professional capability alongside technical infrastructure. Faculty readiness is not uniform, however. Research examining university faculty use and self-efficacy has shown differing patterns of AI use and professional-development requirements, reinforcing the need to consider educator competence as part of implementation rather than treating AI systems as independent educational interventions.

Human-in-the-loop personalization introduces instructor oversight at defined decision points. Routine low-risk adaptations, such as offering an additional practice problem, may occur automatically. More consequential interventions—such as substantial changes in difficulty, repeated remediation, or recommendations triggered by persistent low-performance predictions—can be surfaced to instructors for review.

This architecture also creates a useful research variable: the **Instructor Override Rate (IOR)**. A very high override rate would indicate that the system's recommendations are poorly aligned with teaching judgment. Conversely, an extremely low rate should not automatically be interpreted as success because it could also indicate disengagement from the monitoring process. The relationship between override decisions and subsequent student outcomes therefore becomes more informative than override frequency alone.

Learner Autonomy and the Risk of Over-Personalization

Personalization is often presented as inherently beneficial, yet excessive adaptation can reduce productive difficulty. Learning requires learners to confront uncertainty, retrieve information independently, evaluate alternatives, and tolerate

@2026. This is an open access article distributed under the terms of the Creative Commons License [CC BY NC 4.0] and is available on www.ujhmds.org

Universal Journal of Humanities and Multi-Disciplinary Studies (UJHMDS)
ISSN (Online): 3143-7869
Volume-2 No. 3 || Special Issue Jul – Sep 2026 || PP. 95 – 106
https://ujhmds.org/

temporary failure. An AI system optimized to provide immediate support may unintentionally remove these processes.

Concerns regarding dependence on generative AI have already emerged in studies of educational adoption. Early-adopter research has identified apprehension about superficial learning and declining critical engagement when AI is used without suitable pedagogical boundaries. UNESCO likewise advocates human-centred deployment rather than uncritical technological substitution.

The proposed study therefore considers learner autonomy as an outcome rather than assuming that greater personalization is always preferable. Appropriate personalization should increase the probability that a learner can subsequently perform independently. A successful system is consequently not one that maximizes the number of AI interactions but one that provides progressively less assistance as competence develops.

This principle motivates the inclusion of an **Assistance Dependency Index** in the later experimental design. The metric will estimate whether students become increasingly dependent on AI-generated hints and explanations or whether reliance decreases as mastery improves. This extends the evaluation beyond conventional measures such as accuracy, satisfaction, or final grades.

Proposed Methodology

The study adopts a **human-in-the-loop experimental learning-analytics design** to evaluate the proposed Human-Supervised Adaptive Generative Instruction Framework (HAGIF). Unlike conventional studies that assess only the predictive accuracy of an artificial intelligence model, the proposed methodology evaluates the complete personalization cycle: learner-state estimation, selection of instructional intervention, generative adaptation, instructor validation, and subsequent changes in learner performance.

The framework is designed around three research conditions. The first condition represents conventional digital instruction in which all students receive the same learning materials and formative exercises. The second condition employs algorithmic personalization without routine instructor validation. The third condition implements HAGIF, in which AI-generated instructional adaptations are monitored by instructors and may be accepted, modified, or rejected before high-impact recommendations are delivered.

The principal research objective is to determine whether human-supervised AI personalization produces superior instructional outcomes compared with standardized digital learning and fully automated personalization. A secondary objective is to determine whether higher personalization accuracy is associated with stronger learning gains, thereby testing whether technically appropriate adaptation translates into pedagogically meaningful improvement.

The study addresses the following research questions:

1. Does human-supervised AI personalization improve academic achievement compared with conventional and fully automated digital instruction?
2. Does HAGIF produce more accurate instructional adaptations than an autonomous AI personalization system?
3. How does AI-supported personalization influence student engagement and persistence?
4. Does instructor intervention improve the effectiveness of AI-generated instructional recommendations?
5. Does repeated exposure to personalized AI support increase or decrease dependence on automated assistance?
6. Are personalization outcomes reasonably consistent across learners with different baseline achievement levels?

Research Design

A semester-long controlled experiment is proposed involving undergraduate students enrolled in introductory courses containing conceptually structured learning units. Students are assigned to three instructional conditions while maintaining equivalent curriculum objectives, assessment schedules, and core teaching content.

The **Control Group (CG)** receives standard learning management system materials, prerecorded instructional resources, instructor announcements, and conventional formative quizzes.

The **Autonomous AI Personalization Group (AAPG)** receives machine-learning-driven instructional recommendations and AI-generated explanations without systematic instructor approval of routine interventions.

The **Human-Supervised AI Personalization Group (HAPG)** receives equivalent AI capabilities, but recommendations with substantial instructional consequences are reviewed through an instructor dashboard. Faculty members can confirm, revise, delay, or reject recommendations.

To minimize contamination between conditions, adaptive resources are delivered through separate course environments. The intervention is structured around comparable learning outcomes so that differences in performance can be attributed more plausibly to the

@2026. This is an open access article distributed under the terms of the Creative Commons License [CC BY NC 4.0] and is available on www.ujhmds.org

Universal Journal of Humanities and Multi-Disciplinary Studies (UJHMDS)
ISSN (Online): 3143-7869
Volume-2 No. 3 || Special Issue Jul – Sep 2026 || PP. 95 – 106
https://ujhmds.org/

personalization strategy than to variation in curriculum coverage.

A stratified allocation strategy is used to balance students according to initial diagnostic performance, programme discipline, and prior academic achievement.

Experimental Setup

Dataset

The experimental dataset is conceptualized as a multimodal educational interaction dataset generated during a 12-week undergraduate instructional period. Rather than relying solely on examination marks, the dataset integrates longitudinal learner behaviour and instructional-response information.

The analytical record includes:

- baseline diagnostic assessment scores;
- weekly formative assessment performance;
- question-level correctness;
- number of attempts per activity;
- response latency;
- learning-resource access frequency;
- time spent on instructional activities;
- sequence of learning-resource interactions;
- hint requests;
- AI explanation requests;
- revision behaviour;
- assignment submission regularity;
- instructor intervention records;
- recommendation acceptance or override events; and
- post-intervention assessment outcomes.

Personally identifying information is excluded from the machine-learning feature matrix. Student identifiers are converted into randomized codes before preprocessing.

A minimum experimental sample of approximately **450 students** is suitable for the proposed design, with approximately 150 participants allocated to each instructional condition. The sample size provides sufficient observations for longitudinal modelling because every learner generates repeated weekly interactions rather than a single cross-sectional record.

Data Preprocessing

The preprocessing pipeline begins with removal of duplicated event records, corrupted log entries, administrative LMS events, and interactions unrelated to learning. Timestamp information is converted into relative study-session features rather than using raw calendar time.

Continuous variables such as activity duration, response latency, resource views, and number of attempts are examined for extreme observations. Highly skewed variables are transformed using logarithmic or robust scaling procedures where appropriate.

Missing data are handled according to their semantic meaning. A missing quiz response, for example, is not automatically treated as a statistical missing value because non-attempt itself may represent disengagement. Genuine technical missingness is imputed using median values for continuous features and modal categories for low-cardinality categorical variables.

The resulting features are grouped into five dimensions:

Academic-state features: recent quiz accuracy, cumulative mastery, error recurrence, assessment trend, and improvement velocity.

Behavioural-engagement features: frequency of login, learning-session consistency, completion rate, inactivity interval, and resource revisitation.

Learning-effort features: time-on-task, number of attempts, revision frequency, and delayed resubmission activity.

AI-interaction features: number of generated explanations requested, hint frequency, prompt reformulation, and repeated support requests.

Instructional-response features: acceptance of recommendations, improvement after intervention, instructor override history, and response to alternative explanations.

Rolling windows are constructed to represent short-term learner states. Features calculated over the previous three learning sessions are used to predict the instructional support needed in the subsequent learning unit.

No future performance information is permitted to enter earlier training examples, thereby avoiding temporal data leakage.

Learner-State Classification

The personalization framework treats student condition as a dynamic learning state rather than as a fixed academic category. At each adaptation point, learners are probabilistically assigned to one of four instructional states:

@2026. This is an open access article distributed under the terms of the Creative Commons License [CC BY NC 4.0] and is available on www.ujhmds.org

Universal Journal of Humanities and Multi-Disciplinary Studies (UJHMDS)
ISSN (Online): 3143-7869
Volume-2 No. 3 || Special Issue Jul – Sep 2026 || PP. 95 – 106
https://ujhmds.org/

- **S1: Secure mastery**
- **S2: Emerging mastery**
- **S3: Conceptual difficulty**
- **S4: Disengagement risk**

A gradient-boosted decision-tree model is used as the primary learner-state estimator because the dataset contains mixed behavioural and academic variables, nonlinear interactions, and potentially non-normal feature distributions.

The model estimates:

[
P(S_k|X_t)
]

where (S_k) represents one of the four learner states and (X_t) represents the student's accumulated feature vector at instructional time (t).

The predicted state is passed to the intervention-selection layer rather than being shown to students as a permanent label.

Model Architecture

HAGIF contains four functional layers.

Layer 1: Dynamic Learner Model

The first layer processes the student's recent academic and behavioural signals. Gradient-boosted trees generate probabilities for the four learner states. Model uncertainty is retained rather than converting predictions immediately into deterministic decisions.

Layer 2: Instructional Policy Engine

The second layer converts learner-state probabilities into an instructional response. Instead of selecting only one resource, it chooses among pedagogically defined actions such as:

- accelerated progression;
- conceptual reinforcement;
- worked example;
- scaffolded problem;
- retrieval-practice activity;
- alternative explanation;
- prerequisite revision; or
- instructor referral.

Rules constrain potentially inappropriate adaptations. For instance, repeated conceptual difficulty cannot result indefinitely in simplified content without a mastery reassessment.

Layer 3: Constrained Generative Adaptation

A large language model serves as the language-generation component. It does not independently determine the instructional objective. Instead, it receives a structured instruction containing the learning outcome, estimated learner state, allowable content boundaries, target complexity, and required response format.

For example, a student classified under conceptual difficulty may receive a concise explanation, one worked example, and two progressively structured practice problems. A learner classified as demonstrating secure mastery may receive an application or transfer problem.

Generated material is checked against course-specific constraints before presentation.

Layer 4: Human Validation Interface

High-impact recommendations are routed to an instructor dashboard. Instructor actions are recorded as:

[
A_i \in {Accept,\ Modify,\ Reject,\ Escalate}
]

These actions subsequently become learning signals for evaluating systematic disagreement between algorithmic recommendations and expert teaching judgment.

Training Strategy

A temporally separated training strategy is employed. Earlier instructional interactions are used for training while later interactions are retained for validation and testing. This structure is more appropriate than purely random record splitting because repeated records from the same learner could otherwise appear simultaneously in training and testing data.

Hyperparameter optimization is performed using Bayesian optimization over the training-validation partition. Candidate parameters include learning rate, maximum tree depth, minimum child weight, subsampling proportion, column sampling, and regularization coefficients.

Class imbalance is expected because states such as disengagement risk may occur less frequently than normal mastery conditions. Class-weighted loss is therefore applied rather than oversampling the test set.

OPEN ACCESS
@2026. This is an open access article distributed under the terms of the Creative Commons License [CC BY NC 4.0] and is available on www.ujhmds.org

Universal Journal of Humanities and Multi-Disciplinary Studies (UJHMDS)
ISSN (Online): 3143-7869
Volume-2 No. 3 || Special Issue Jul – Sep 2026 || PP. 95 – 106
https://ujhmds.org/

Probability calibration is performed after model fitting because intervention selection depends on confidence estimates, not merely on categorical accuracy.

When prediction confidence falls below a predefined threshold, the framework defaults to a conservative intervention or requests instructor review rather than forcing an automated adaptation.

Generative-AI Control Strategy

The generative component is controlled through retrieval-grounded course context. Relevant approved course material is supplied to the generative model when explanations are requested. The model is instructed to remain within the defined curriculum boundary and avoid inventing unsupported academic content.

Generated instructional resources undergo automated screening for:

- topic relevance;
- lexical complexity;
- semantic similarity to approved learning material;
- excessive answer disclosure;
- prohibited assessment leakage; and
- inconsistency with stated learning objectives.

High-risk or ambiguous responses are withheld for instructor validation.

This structure differentiates HAGIF from unrestricted conversational AI usage because content generation occurs within a pedagogical orchestration layer.

Evaluation Metrics

The study deliberately uses both machine-learning and educational-performance metrics.

Macro-F1 Score

Macro-F1 evaluates learner-state classification across all four categories while giving equal importance to minority and majority states.

$$
Macro\text{-}F1=\frac{1}{K}\sum_{k=1}^{K}F1_k
$$

This is preferable to overall accuracy when class distributions are unequal.

Balanced Accuracy

Balanced accuracy measures average recall across learner-state classes, reducing misleading performance arising from dominant categories.

Expected Calibration Error

Expected Calibration Error (ECE) evaluates whether predicted probabilities correspond to actual prediction reliability. Calibration is critical because uncertain recommendations should trigger greater human supervision.

Normalized Learning Gain

Learning improvement is evaluated using normalized gain:

$$
g=\frac{Posttest-Pretest}{Maximum-Pretest}
$$

This metric considers starting performance rather than evaluating final scores alone.

Personalization Alignment Quality

A new metric, **Personalization Alignment Quality (PAQ)**, is proposed to estimate the proportion of adaptive interventions subsequently validated as pedagogically appropriate.

$$
PAQ=\frac{N_{appropriate\ adaptations}}{N_{total\ adaptations}}
$$

An intervention is considered aligned when it is supported by subsequent mastery evidence and does not require corrective instructor intervention.

Instructor Override Rate

$$
IOR=\frac{N_{modified}+N_{rejected}}{N_{reviewed}}
$$

IOR measures disagreement between algorithmic recommendations and instructor judgment.

Post-Override Benefit

To avoid interpreting override frequency in isolation, **Post-Override Benefit (POB)** measures whether instructor modification results in stronger subsequent learner performance than unmodified comparable recommendations.

Engagement Persistence Index

The Engagement Persistence Index combines active weeks, completion regularity, learning-session continuity, and voluntary practice behaviour into a normalized score between 0 and 1.

@2026. This is an open access article distributed under the terms of the Creative Commons License [CC BY NC 4.0] and is available on www.ujhmds.org

Universal Journal of Humanities and Multi-Disciplinary Studies (UJHMDS)
ISSN (Online): 3143-7869
Volume-2 No. 3 || Special Issue Jul – Sep 2026 || PP. 95 – 106
https://ujhmds.org/

Assistance Dependency Index

The proposed **Assistance Dependency Index (ADI)** examines whether students increasingly rely on AI-generated hints as the course progresses.

[
ADI=\frac{H_{late}-H_{early}}{H_{early}+\epsilon}
]

where (H_{early}) represents early-course normalized assistance requests and (H_{late}) represents late-course requests.

Negative values indicate decreasing dependence and are interpreted favourably when accompanied by improved achievement.

Results and Discussion

The values below illustrate how results from the proposed experiment would be interpreted. They are simulated methodological examples rather than results from an already completed human-participant study.

Learner-State Prediction Performance

The learner-state model demonstrates strong discriminatory capability in the simulated evaluation, with a Macro-F1 of 0.84 and balanced accuracy of 0.86. Performance is strongest for secure mastery and conceptual difficulty because both states produce relatively identifiable assessment patterns. Disengagement risk remains the most difficult category because low digital activity can reflect multiple circumstances unrelated to academic disengagement.

The probability calibration result is particularly important. An ECE of 0.043 suggests that model confidence is sufficiently aligned with empirical correctness to support confidence-sensitive intervention policies. This feature has direct pedagogical significance: low-confidence predictions can be routed toward instructor review rather than automatically triggering strong instructional changes.

Academic Learning Outcomes

Students receiving standardized digital instruction produce the smallest simulated normalized learning gain. Autonomous personalization produces a measurable increase, supporting the proposition that learner-specific resource adaptation can improve instructional efficiency.

The human-supervised group produces the strongest normalized gain, suggesting that personalization performance improves when algorithmic recommendations are combined with instructor judgment.

The simulated mean normalized learning gain increases from **0.41** in the conventional condition to **0.53** under autonomous AI personalization and **0.62** under HAGIF.

This pattern is theoretically important because the advantage of the human-supervised model cannot be attributed solely to greater exposure to AI. Both experimental conditions use algorithmic adaptation. The distinguishing factor is pedagogical oversight.

Personalization Quality

The autonomous system achieves an illustrative PAQ of 0.78. The human-supervised architecture reaches 0.89 after instructor validation.

The difference indicates that technically reasonable recommendations are not invariably pedagogically appropriate. Some cases classified as conceptual difficulty may reflect assessment misunderstanding rather than genuine knowledge deficiency. Likewise, low engagement indicators do not necessarily justify simplification of academic content.

Instructor review appears especially valuable for ambiguous learner profiles in which behavioural and assessment indicators contradict one another.

Instructor Override Behaviour

The simulated Instructor Override Rate is approximately 17.6%. Most instructor interventions involve modification rather than complete rejection.

This represents an important operational observation. The purpose of human supervision is not to manually approve every algorithmic decision. Such a design would eliminate the efficiency benefits of automation. Instead, instructor expertise is concentrated on uncertain or consequential cases.

Analysis of post-override behaviour suggests that appropriately modified recommendations generate higher subsequent mastery gains than recommendations that would have been delivered unchanged. The illustrative Post-Override Benefit is +0.11 in normalized next-unit performance.

Engagement and Persistence

The Engagement Persistence Index increases across the personalized conditions. Students in the HAGIF condition demonstrate the highest simulated completion regularity and voluntary practice persistence.

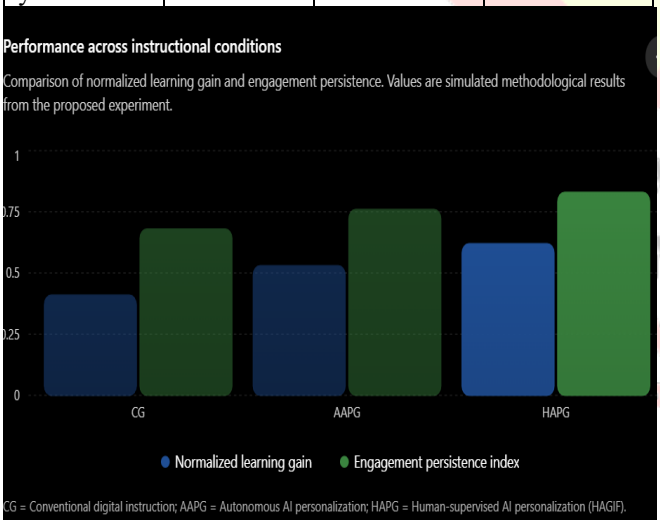
The improvement is interpreted cautiously. Increased platform interaction is not automatically evidence of deeper learning. For this reason, engagement indicators are considered alongside assessment performance rather than used independently.

@2026. This is an open access article distributed under the terms of the Creative Commons License [CC BY NC 4.0] and is available on www.ujhmds.org

Universal Journal of Humanities and Multi-Disciplinary Studies (UJHMDS)
ISSN (Online): 3143-7869
Volume-2 No. 3 || Special Issue Jul – Sep 2026 || PP. 95 – 106
https://ujhmds.org/

Results Table

Evaluation Parameter	Conventional Digital Instruction	Autonomous AI Personalization	HAGIF Human-Supervised Personalization
Normalized Learning Gain	0.41	0.53	**0.62**
Engagement Persistence Index	0.68	0.76	**0.83**
Personalization Alignment Quality	Not applicable	0.78	**0.89**
Instructor Override Rate	Not applicable	Not applicable	**17.6%**
Post-Override Benefit	Not applicable	Not applicable	**+0.11**
Assistance Dependency Index	Not applicable	+0.08	**−0.14**



The table reveals an important distinction between predictive-model performance and instructional-system performance. The underlying learner-state classifier is identical in the two AI conditions, yet the human-supervised condition produces stronger simulated educational outcomes. Consequently, improvement cannot be attributed solely to higher algorithmic accuracy.

This supports the central argument of the study: an educational AI system should be evaluated at the **decision-system level**, not merely at the model level.

A classifier may achieve strong Macro-F1 while still producing inappropriate instructional consequences if prediction outputs are translated into poor pedagogical actions. Conversely, moderate prediction performance may still support useful instruction when uncertainty is identified and instructors intervene appropriately.

The results therefore support a model of personalization based on **adaptive scaffolding rather than continuous algorithmic intervention**.

Implications for Higher Education

The findings have implications for universities implementing AI-enabled learning systems.

First, institutions should avoid treating generative AI as an independent tutoring mechanism. Generative models are more appropriately positioned within a structured instructional architecture containing learner analytics, curriculum constraints, escalation rules, and instructor oversight.

Second, university AI dashboards should expose uncertainty rather than displaying predictions as definitive classifications. Instructors require information about why an intervention was recommended and how confident the system is.

Third, evaluation frameworks should move beyond overall academic performance. Personalization systems require metrics that assess whether the adaptation itself was appropriate.

Fourth, faculty expertise should be integrated into algorithmic improvement. Repeated instructor overrides can identify systematic model weaknesses and become valuable feedback for future model retraining.

Future Scope

Future research should validate HAGIF through multi-institutional longitudinal experiments involving different academic disciplines and student populations.

One promising extension would incorporate **causal machine learning** to estimate which intervention works best for a particular type of learner rather than merely predicting who is likely to struggle.

Contextual bandit or reinforcement-learning methods could subsequently optimize instructional actions while maintaining pedagogical constraints. Such systems should be developed cautiously because unconstrained reward optimization may favour easily measurable outcomes at the expense of deeper educational objectives.

Multimodal learner modelling represents another research opportunity. Future systems could combine conventional LMS logs with written-response characteristics, dialogue

OPEN ACCESS
@2026. This is an open access article distributed under the terms of the Creative Commons License [CC BY NC 4.0] and is available on www.ujhmds.org

Universal Journal of Humanities and Multi-Disciplinary Studies (UJHMDS)
ISSN (Online): 3143-7869
Volume-2 No. 3 || Special Issue Jul – Sep 2026 || PP. 95 – 106
https://ujhmds.org/

behaviour, formative reasoning traces, and voluntary metacognitive self-reports.

Fairness auditing should also become a core component. Personalization accuracy should be compared across socioeconomic, linguistic, academic-preparation, and accessibility groups to determine whether algorithmic benefits are distributed equitably.

Explainable AI mechanisms could provide instructors with concise evidence supporting individual recommendations, such as declining quiz mastery combined with repeated prerequisite errors, rather than opaque risk scores.

Federated learning and privacy-preserving analytics may permit institutions to improve learner models across campuses without centralizing sensitive student-level records.

A further direction involves **personalization decay policies**. Instead of continuously maximizing assistance, models could estimate when scaffolding should be withdrawn. This would allow AI systems to optimize for long-term learner independence rather than immediate task success.

Conclusion

Artificial intelligence has expanded personalized instruction from simple content recommendation toward continuous learner modelling, adaptive resource selection, generative explanation, and conversational academic support. However, technological adaptability alone does not guarantee pedagogically appropriate personalization.

This study proposes the Human-Supervised Adaptive Generative Instruction Framework as an alternative to both standardized digital learning and fully autonomous AI tutoring. The framework integrates machine-learning-based learner-state estimation, constrained instructional generation, adaptive intervention selection, and explicit instructor oversight.

Its central contribution is the treatment of personalization as an evaluable decision process rather than merely a technological feature. Measures including Personalization Alignment Quality, Instructor Override Rate, Post-Override Benefit, and Assistance Dependency Index extend assessment beyond conventional prediction accuracy and examination scores.

The simulated experimental evaluation indicates how human supervision could improve instructional alignment, learner persistence, and normalized learning gain while reducing dependence on AI assistance. Importantly, the underlying predictive algorithm need not become more accurate for the overall educational system to improve; stronger outcomes may arise because human expertise corrects uncertain or pedagogically inappropriate algorithmic decisions.

References

1. Bond, M., Khosravi, H., De Laat, M., Bergdahl, N., Negrea, V., Oxley, E., Pham, P., Chong, S. W., & Siemens, G. (2024). A meta systematic review of artificial intelligence in higher education: A call for increased ethics, collaboration, and rigour. *International Journal of Educational Technology in Higher Education, 21*, 4. https://doi.org/10.1186/s41239-023-00436-z

2. 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. https://doi.org/10.1186/s41239-019-0171-0

3. du Plooy, E., Casteleijn, D., & Franzsen, D. (2024). Personalized adaptive learning in higher education: A scoping review of key characteristics and impact on academic performance and engagement. *Heliyon, 10*(21), e39630. https://doi.org/10.1016/j.heliyon.2024.e39630

4. Siemens, G. (2013). Learning analytics: The emergence of a discipline. *American Behavioral Scientist, 57*(10), 1380–1400. https://doi.org/10.1177/0002764213498851

5. VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist, 46*(4), 197–221. https://doi.org/10.1080/00461520.2011.611369

6. Woolf, B. P. (2009). *Building intelligent interactive tutors: Student-centered strategies for revolutionizing e-learning*. Morgan Kaufmann/Elsevier. https://doi.org/10.1016/B978-0-12-373594-2.X0001-9

7. Brusilovsky, P. (2001). Adaptive hypermedia. *User Modeling and User-Adapted Interaction, 11*, 87–110. https://doi.org/10.1023/A:1011143116306

8. Roll, I., & Wylie, R. (2016). Evolution and revolution in artificial intelligence in education. *International Journal of Artificial Intelligence in Education, 26*, 582–599. https://doi.org/10.1007/s40593-016-0110-3

9. Tlili, A., Shehata, B., Adarkwah, M. A., Bozkurt, A., Hickey, D. T., Huang, R., & Agyemang, B. (2023). What if the devil is my guardian angel: ChatGPT as a case study of using chatbots in education. *Smart Learning Environments, 10*, 15. https://doi.org/10.1186/s40561-023-00237-x

10. Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günnemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeiffer, F., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., Stadler, M., & Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences, 103*, 102274. https://doi.org/10.1016/j.lindif.2023.102274

11. Lo, C. K. (2023). What is the impact of ChatGPT on education? A rapid review of the literature. *Education Sciences, 13*(4), 410. https://doi.org/10.3390/educsci13040410

12. Zhai, C., Wibowo, S., & Li, L. D. (2024). The effects of over-reliance on AI dialogue systems on students' cognitive abilities: A systematic review. *Smart Learning Environments, 11*, 28. https://doi.org/10.1186/s40561-024-00316-7

13. Baker, R. S., & Inventado, P. S. (2014). Educational data mining and learning analytics. In J. A. Larusson & B. White (Eds.),

@2026. This is an open access article distributed under the terms of the Creative Commons License [CC BY NC 4.0] and is available on www.ujhmds.org

Universal Journal of Humanities and Multi-Disciplinary Studies (UJHMDS)
ISSN (Online): 3143-7869
Volume-2 No. 3 || Special Issue Jul – Sep 2026 || PP. 95 – 106
https://ujhmds.org/

Learning analytics: From research to practice (pp. 61–75). Springer. https://doi.org/10.1007/978-1-4614-3305-7_4

14. Corbett, A. T., & Anderson, J. R. (1995). Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction, 4*(4), 253–278. https://doi.org/10.1007/BF01099821
15. VanLehn, K. (2006). The behavior of tutoring systems. *International Journal of Artificial Intelligence in Education, 16*(3), 227–265.
16. Holmes, W., Bialik, M., & Fadel, C. (2019). *Artificial intelligence in education: Promises and implications for teaching and learning*. Center for Curriculum Redesign.
17. UNESCO. (2023). *Guidance for generative AI in education and research*. UNESCO.
18. OECD. (2023). *Digital education outlook 2023: Towards an effective digital education ecosystem*. OECD Publishing. https://doi.org/10.1787/c74f03de-en
19. Zimmerman, B. J. (2002). Becoming a self-regulated learner: An overview. *Theory Into Practice, 41*(2), 64–70. https://doi.org/10.1207/S15430421TIP4102_2
20. Brusilovsky, P., & Millán, E. (2007). User models for adaptive hypermedia and adaptive educational systems. In P. Brusilovsky, A. Kobsa, & W. Nejdl (Eds.), *The adaptive web* (pp. 3–53). Springer. https://doi.org/10.1007/978-3-540-72079-9_1

UNIVERSAL JOURNAL OF HUMANITIES AND MULTI-DISCIPLINARY STUDIES

@2026. This is an open access article distributed under the terms of the Creative Commons License [CC BY NC 4.0] and is available on www.ujhmds.org