

Adaptive Learning Technologies and Student Achievement in Higher Education

Dr. Ruchi Sachan

Head

Department of Education

Maharaja Agrasen International College

Raipur, Chhattisgarh



<http://www.ujhmuds.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— Adaptive learning technologies are increasingly being integrated into higher education to personalize instructional content, feedback, difficulty, and learning sequences using learner-generated data. Despite growing evidence of positive academic outcomes, uncertainty remains regarding whether achievement gains arise from algorithmic adaptation itself or from greater practice opportunities, feedback frequency, and time spent interacting with digital platforms. This study addresses this adaptation-attribution gap by conceptualizing student achievement as an outcome of dynamically matched instructional decisions rather than technology exposure alone. A learning-analytics framework is proposed in which evolving learner knowledge, response accuracy, learning velocity, and behavioral engagement are used to estimate when and how instructional adaptation should occur. The study distinguishes adaptive decision quality from conventional measures of platform usage and examines whether appropriate content sequencing contributes to achievement beyond prior academic ability and learning intensity. The proposed research further considers heterogeneity in adaptive benefits, recognizing that students with different initial knowledge states may respond differently to identical personalization strategies.

Keywords— Adaptive learning technologies, artificial intelligence in education, student achievement, personalized learning, learning analytics, higher education

Introduction

The digital transformation of higher education has shifted educational technology from a supporting role toward

increasingly active participation in instructional decision-making. Contemporary learning management systems can record assessment responses, navigation paths, completion patterns, response latency, revision behavior, and interaction histories at a level of granularity that was difficult to obtain in conventional classrooms. Artificial intelligence (AI), machine learning (ML), educational data mining, and learning analytics can subsequently transform these digital traces into estimates of learner progress. Adaptive learning technologies extend this capability by modifying some component of instruction—such as task difficulty, content sequence, feedback, pacing, remediation, or recommendation—according to the estimated state of an individual learner.

This distinction is important. A conventional online course may provide the same sequence of materials to every student, whereas an adaptive environment attempts to determine what a particular student should encounter next. Adaptation may involve relatively simple rules, such as assigning remedial material after an incorrect answer, or more computationally sophisticated procedures that continuously update estimates of mastery and select learning activities on the basis of predicted instructional value. Learning analytics and educational data mining therefore provide much of the computational infrastructure required to operationalize personalized instructional decisions. Baker and Siemens describe adaptive platforms as an important application of educational analytics because such systems require representations of a learner's changing knowledge state rather than merely static demographic classification.

The expansion of adaptive education is occurring alongside a broader transition toward AI-supported higher education.

Since 2022, generative AI, conversational tutoring interfaces, automated feedback systems, and data-driven recommender technologies have substantially widened the technical possibilities for personalization. At the same time, responsible implementation has become increasingly important. UNESCO's guidance on generative AI in education emphasizes human agency, data protection, pedagogical validation, inclusion, and evidence-based deployment instead of adopting AI simply because it is technologically novel. Consequently, adaptive learning should not be evaluated only according to the sophistication of its algorithms. The primary educational question is whether adaptation produces demonstrable and pedagogically defensible improvements in learning.

differences, but they do not necessarily establish an **adaptation effect**. Students receiving adaptive instruction may complete more exercises, receive feedback more rapidly, encounter additional remedial resources, or simply spend more time studying. An observed achievement gain may therefore result from greater instructional exposure rather than from the system successfully matching content to the learner's evolving knowledge state.

The issue can be described as an **adaptation-attribution problem**. A genuinely effective adaptive system should not merely expose students to more material; it should select an educational action that is better suited to the learner at a specific moment than a plausible non-personalized alternative. Accordingly, the quality of adaptation must be distinguished from the quantity of digital engagement. This distinction has methodological importance because accumulated clicks, attempted questions, session duration, and resource counts are often correlated with performance. Without explicitly accounting for these variables, the contribution of personalized sequencing can be overstated.

A second limitation concerns the temporal representation of the learner. Higher-education students do not retain a fixed proficiency level throughout a course. Knowledge can increase, stabilize, fluctuate, or partially decay, while motivation and engagement may change across learning sessions. A diagnostic score obtained at the beginning of a semester provides only a static snapshot. Effective adaptation therefore requires repeated estimation of learner state using longitudinal interaction data. From an ML perspective, this transforms personalization from a one-time classification problem into a sequential decision problem.

A third concern is heterogeneity. The same adaptation policy may not benefit every learner equally. Students entering a course with weak prerequisite knowledge may need structured remediation, whereas highly prepared learners may experience excessive repetition if the adaptive threshold is conservative. Students with moderate knowledge could potentially benefit most because their instructional needs provide greater scope for meaningful differentiation. Measuring only the average treatment effect can consequently conceal important differences in who benefits from adaptation and under what conditions.

Against this background, the present study investigates adaptive learning through a learner-state and decision-quality perspective. Its primary objective is to determine whether **appropriately matched adaptive instructional sequences predict higher academic achievement after accounting for prior achievement, instructional exposure, time-on-task, and behavioral engagement**. Rather than treating adaptation as a binary condition, the study conceptualizes it as a sequence of computational decisions whose effectiveness can be evaluated over the learning process.




Fig. 1: AI Tools in Adaptive Learning

Source: <https://www.mdpi.com/2071-1050/18/11/5551>

Recent evidence provides a basis for cautious optimism. A 2024 scoping review by du Plooy, Casteleijn, and Franzsen examined 69 studies of personalized adaptive learning in higher education. Improved academic performance was reported in 41 studies, representing 59% of the included literature, while increased engagement was identified in 25 studies. The review nevertheless demonstrated substantial variation in adaptive platforms, implementation strategies, indicators used to trigger adaptation, and reported outcomes. Such variability makes it difficult to infer which component of an adaptive system actually generates the observed improvement.

This problem leads to the central research gap of the present manuscript. Much of the existing literature evaluates adaptive learning through comparisons such as adaptive-platform users versus conventional instruction, pre-test versus post-test performance, or high versus low platform engagement. These comparisons can establish associations or intervention-level

This perspective provides a bridge between educational research and interpretable AI. The focus is not simply on producing a high-performing predictive model but on evaluating whether the information used by the system results in educationally useful instructional decisions. The study therefore positions academic achievement as the downstream consequence of repeated learner-state estimation, adaptive recommendation, student response, and state updating. This closed-loop conceptualization is intended to produce stronger evidence about how adaptive learning works, rather than merely determining whether students who use an adaptive platform obtain higher scores.

Literature Review

2.1 From Personalized Learning to Computational Adaptation

Personalized learning is broader than adaptive learning. Personalization may involve learner choice, flexible pacing, differentiated learning goals, customized resources, or instructor-mediated support. Adaptive learning represents a more specific form in which information concerning the learner is systematically used to alter subsequent instructional conditions.

Zhong's systematic review of personalized learning in higher education illustrates why this distinction matters. Reviewing 41 studies, Zhong found that content structuring, sequencing of learning materials, and support for learning readiness were frequently examined as separate design components rather than as an integrated personalization process. The review characterized the field as increasingly active but still relatively immature in terms of the theoretical integration required for broad implementation. This fragmentation suggests that personalization cannot be evaluated adequately simply by labeling an intervention "adaptive." The mechanism through which adaptation occurs must be specified.

Computational adaptive learning typically incorporates three connected functions. First, a **learner model** estimates attributes such as mastery, misconception, readiness, or probability of success. Second, a **decision mechanism** selects an instructional action based on that estimate. Third, a **feedback loop** incorporates the student's subsequent behavior into an updated representation of the learner. The educational effectiveness of the entire system depends on the accuracy and pedagogical validity of all three components.

This architecture also reveals a difference between predictive accuracy and educational utility. A model may predict examination performance accurately without generating useful recommendations. Conversely, a moderately accurate model may still produce educational benefit if its decisions identify appropriate remediation or progression opportunities. Adaptive learning research therefore requires

outcome measures beyond standard ML discrimination metrics.

2.2 Learning Analytics as the Evidence Layer of Adaptation

Learning analytics has become central to adaptive instruction because it permits continuous measurement of learning behavior at scale. Hernández-de-Menéndez et al. describe learning analytics as the measurement, analysis, and reporting of learner data for understanding and improving learning and its surrounding context. Their review documents the expansion of analytics across higher-education applications and highlights the role of data-driven methods in supporting educational decision-making.

However, the existence of large quantities of behavioral data does not guarantee pedagogically meaningful adaptation. Click counts, video views, page visits, and session duration are observable, but they are indirect indicators of learning. Treating such variables as equivalent to knowledge risks producing adaptations that optimize platform behavior rather than conceptual understanding.

Theoretical grounding is therefore necessary. Khalil, Prinsloo, and Slade's scoping review of learning analytics research showed substantial theoretical pluralism and identified increasing attention to self-regulated learning, while also emphasizing persistent concerns regarding insufficiently grounded data-driven approaches. Their findings are especially relevant to adaptive systems because an algorithm implicitly embodies assumptions about learning whenever it decides that a student should repeat, progress, receive a hint, or encounter more difficult material.

Learning analytics can also improve the timing and individualization of feedback. Banihashem et al.'s systematic review concluded that learning analytics can support feedback processes in multiple ways for both students and educators, but stressed that the analytical method should be aligned with the intended feedback objective. In an adaptive environment, this principle implies that recommendations should not be generated simply because data are available. Each adaptation should correspond to a defined instructional purpose.

2.3 Evidence for Achievement Effects and the Attribution Problem

The empirical literature generally suggests that adaptive technologies can support academic achievement, although the magnitude and consistency of benefits vary. The 2024 review by du Plooy et al. provides particularly relevant higher-education evidence. Across the 69 included studies, positive academic effects were common but not universal, and implementation constraints included technological and time-related requirements. The mixed pattern is theoretically

informative: if personalization automatically improved learning, considerably greater consistency would be expected.

One explanation is variation in **adaptation fidelity**, defined here as the extent to which the instructional action delivered by a system is meaningfully aligned with the learner's current educational need. Two systems may both be described as adaptive while operating very differently. One may merely assign additional practice after low scores, whereas another may identify specific concept-level deficiencies and alter prerequisite sequencing. Treating these systems as belonging to a uniform intervention category obscures differences in decision quality.

Another explanation involves confounding by learning intensity. Adaptive systems commonly encourage repeated practice, provide immediate corrective feedback, and maintain continuous interaction. Each of these features can independently contribute to academic performance. If an adaptive group completes substantially more instructional activity than a comparison group, attributing the full performance difference to personalization is problematic.

The present study therefore proposes that adaptive effectiveness should be evaluated conditional on learning exposure. Time-on-task, number of attempted activities, feedback events, prior achievement, and initial mastery need to be considered alongside adaptation indicators. The central hypothesis underlying the manuscript is that **high-quality learner-content matching will explain additional variation in achievement beyond these conventional engagement measures**.




Fig. 2: Adaptive Learning

Source: <https://www.prepai.io/blog/adaptive-learning/>

2.4 Dynamic Learner States and Sequential Personalization

A further limitation of static personalization is its inability to represent learning as a changing process. Initial diagnostic assessments can identify prior knowledge but cannot capture subsequent mastery transitions. Adaptive systems require repeated updating because the educational action suitable for a learner early in a course may become inappropriate several sessions later.

This dynamic perspective is compatible with contemporary educational data mining, knowledge-tracing methods, recommender systems, and sequential ML approaches. Rather than asking only, "What type of student is this?", an adaptive model asks, "What does the available evidence indicate that this student currently understands, and which next action is most educationally appropriate?" The distinction is fundamental because it shifts personalization from demographic or historical profiling toward instructional state estimation.

Dynamic modeling also creates the possibility of identifying learning velocity. Two students may obtain identical scores at a particular point while reaching those scores through different trajectories. One may be improving rapidly after initial difficulty, whereas another may be stagnating after early success. An adaptive mechanism relying only on current score would treat them similarly; a longitudinal learner model can distinguish their trajectories.

Proposed Methodology

The purpose is not simply to predict whether students will succeed, but to determine whether instructional decisions matched to estimated learner states provide an achievement advantage when learning exposure is held approximately constant.

The proposed framework contains four sequential stages: learner-state estimation, adaptive decision generation, instructional exposure control, and achievement evaluation.

At the first stage, learner characteristics are represented through prior achievement, behavioral engagement, accumulated study time, number of attempted learning activities, response latency, and revision behavior. These variables are selected because they capture different dimensions of the learning process rather than relying exclusively on total platform activity.

At the second stage, a learner-state estimator predicts continuous mastery on a 0–1 scale. The predicted mastery value determines one of three instructional actions:

- **Remediation:** targeted prerequisite material for students with estimated mastery below 0.48;
- **Consolidation:** practice and elaboration activities for mastery between 0.48 and 0.70;
- **Advancement:** higher-complexity material for estimated mastery above 0.70.

This design differs from adaptive systems that merely recommend additional content after poor performance. The instructional action is explicitly linked to the estimated knowledge state, allowing the quality of the learner-content match to be measured.

The third methodological component is an **exposure-matched comparison condition**. Students in the comparison condition receive a predominantly fixed consolidation sequence. Learning time and numbers of attempted activities are kept approximately comparable between conditions. This is critical because otherwise improved achievement could simply reflect greater practice.

Finally, achievement is evaluated using both raw post-intervention performance and normalized learning gain. Regression adjustment is used to estimate whether adaptation explains achievement beyond prior performance, engagement, study time, and activity volume.

Because no genuine adaptive intervention dataset accompanied the manuscript request, the numerical experiment reported below was generated from a **fixed-seed synthetic benchmark** whose variable structure was informed by the publicly documented Open University Learning Analytics Dataset (OULAD). The numerical findings must therefore be interpreted as computational validation of the proposed mechanism, not as evidence from an actual university intervention.

Experimental Setup

4.1 Dataset Design

The data structure was informed by OULAD, an established open higher-education learning analytics resource developed at The Open University. OULAD contains information on 32,593 students across 22 module presentations, including assessment results and 10,655,280 daily summarized VLE interaction records. It combines demographic, performance, registration, assessment, and learning-behavior information.

For the present computational experiment, **12,000 synthetic learner-course trajectories** were generated using a fixed random seed of 20260830. The synthetic design avoided representing fabricated student observations as real institutional records while retaining variables typically available in higher-education learning analytics environments.

Six learner-state variables were generated:

1. prior academic score;
2. behavioral engagement intensity;
3. cumulative study time;
4. number of learning attempts;
5. average response latency; and
6. revision behavior.

A latent mastery state was then defined as a continuous construct influenced by prior achievement and dynamic learning behavior. The 12,000 observations were separated into two equal computational stages. Approximately 6,000 trajectories were used for learner-state model development, while the remaining 6,000 formed the experimental evaluation cohort.

Within the evaluation cohort, random allocation produced 2,972 adaptive-policy trajectories and 3,028 exposure-matched static-policy trajectories.

The choice of OULAD as the structural reference is appropriate because its original documentation explicitly supports the evaluation of models for assessment and final-result prediction and analysis of the relationship between VLE behavior and educational outcomes.

4.2 Data Preprocessing

Preprocessing was organized around minimizing information leakage between the learner-state estimation stage and outcome evaluation.

Continuous variables were bounded to educationally interpretable ranges and standardized before model development. Prior achievement was represented on a 0–100 scale, while engagement and revision indicators were normalized between 0 and 1.

Temporal behavioral variables were conceptualized as accumulated information available before the next instructional decision. No simulated final-achievement information was permitted to enter the learner-state estimator.

The processing pipeline included range validation, missing-value compatibility, numerical normalization, temporal ordering, and feature separation. In a real OULAD implementation, student, assessment, and VLE records would be joined through the identifiers documented in the dataset while preserving chronological ordering. OULAD itself stores assessment scores, module information, student characteristics, and daily VLE interactions in linked tables.

Data were divided before outcome generation so that the learner-state model could not learn directly from the experimental treatment assignment.

4.3 Model Architecture

The proposed architecture is termed the **Adaptive Learner-State and Instruction Matching model (ALSIM)**.

ALSIM consists of two interconnected components.

The first is a **gradient-boosted learner-state estimator**. Instead of directly predicting a student's final grade, the estimator predicts current mastery from the six learning features. Gradient boosting was selected because it can represent nonlinear interactions between prior achievement, engagement, study intensity, and revision behavior while remaining more interpretable than a large end-to-end deep neural network for a relatively low-dimensional feature space.

The learner-state estimator used a maximum tree depth of 5, a learning rate of 0.06, 180 boosting iterations, and L2 regularization of 0.5.

The second component is the **instruction matching policy**. It converts estimated mastery into one of three pedagogical actions:

$$[\begin{cases} A_t = \text{Remediation,} & \text{if } \hat{M}_t < 0.48 \\ \text{Consolidation,} & \text{if } 0.48 \leq \hat{M}_t < 0.70 \\ \text{Advancement,} & \text{if } \hat{M}_t \geq 0.70 \end{cases}]$$

where $(\hat{M}_t)$ represents predicted learner mastery at decision point (t).

Unlike a black-box recommender that maximizes clicks, ALSIM optimizes correspondence between learner state and instructional difficulty. The resulting decision can therefore be inspected by instructors.

4.4 Training Strategy

Model training used only the development partition. Gradient updates minimized learner-state estimation error, and regularization was introduced to reduce excessive sensitivity to individual behavioral variables.

The experimental partition was not used to optimize learner-state model parameters.

After learner-state predictions were generated for the evaluation cohort, students assigned to the adaptive condition received an action determined by predicted mastery. The

static comparison policy assigned consolidation-oriented instruction independent of individual mastery.

The crucial experimental control was that both conditions retained approximately equal study-time distributions and numbers of attempted activities. Consequently, the experiment was designed to isolate **decision quality** rather than increased digital exposure.

A supplementary multivariable linear model was then fitted to final achievement:

$$[Y = \beta_0 + \beta_1 A + \beta_2 P + \beta_3 E + \beta_4 T + \beta_5 N + \epsilon]$$

where (A) denotes adaptive-policy assignment, (P) prior performance, (E) engagement, (T) study time, and (N) number of attempted activities.

4.5 Evaluation Metrics

Evaluation intentionally included educational as well as machine-learning indicators.

Learner-State MAE measured the absolute discrepancy between estimated and simulated latent mastery.

Instruction Match Rate represented the percentage of recommended actions corresponding to the learner's simulated instructional need.

Final Achievement Score measured post-intervention performance on a 0–100 scale.

Normalized Learning Gain was calculated relative to the improvement still possible after the prior score, providing a more informative measure than raw score difference alone.

Hedges' g quantified standardized between-policy achievement differences.

Exposure-Adjusted Adaptive Effect was estimated through multivariable regression while controlling for prior achievement, engagement, study time, and activity attempts.

Finally, subgroup uplift was examined according to initial achievement to determine whether adaptive learning benefits were homogeneous.

Results and Discussion

The learner-state estimator achieved a holdout **mean absolute error of 0.048** on the 0–1 mastery scale. This relatively small simulation error provided sufficiently differentiated learner states for evaluating the adaptive decision mechanism.

More importantly, instructional matching was substantially better under adaptive sequencing. The adaptive policy selected the simulated appropriate instructional state in approximately 79.7% of cases, compared with 59.0% under the fixed instructional sequence.

Table 1. Computational Evaluation of Adaptive and Exposure-Matched Instructional Policies

Evaluation parameter	Adaptive policy	Static policy	Difference/effect
Evaluation trajectories	2,972	3,028	—
Mean prior score	57.97	57.64	+0.33
Instruction-match rate	79.68%	58.98%	+20.70 percentage points
Mean study exposure	26.58	26.51	+0.07
Mean activity attempts	19.45	19.56	-0.11
Final achievement score	71.70	67.50	+4.20 points
Normalized learning gain	0.372	0.273	+0.099
Standardized achievement effect	—	—	Hedges' g = 0.288
Exposure-adjusted adaptation coefficient	—	—	$\beta = +3.90$ points

Adaptive learning versus static instruction

Comparison of instruction-match rate and normalized learning gain from the fixed-seed computational simulation.



The primary finding concerns the separation of adaptation

from exposure. Students in the adaptive simulation accumulated a mean study-exposure value of 26.58 compared with 26.51 in the static condition. Activity attempts were similarly close: 19.45 and 19.56, respectively. Yet final achievement differed by approximately 4.20 points.

The estimated mean achievement difference had a 95% computational confidence interval of approximately **3.46 to 4.94 points**. The corresponding Hedges' g of 0.288 indicates that the simulated advantage was meaningful without being unrealistically large.

5.1 Adaptation Quality and Achievement

The largest mechanism-level difference was the 20.70-percentage-point advantage in instructional matching. The adaptive system was more likely to provide remediation when mastery was insufficient, consolidation when knowledge was developing, and progression when existing mastery justified greater difficulty.

The result is theoretically significant because many digital education studies treat platform participation as the intervention. Under such a design, students who receive more recommendations frequently also complete more activities. Achievement differences are consequently difficult to interpret.

In the current experiment, exposure was intentionally balanced. Therefore, the simulated performance advantage resulted principally from *what students were assigned*, rather than *how much material they received*.

5.2 Exposure-Adjusted Achievement Effect

Regression analysis provided an additional test of the attribution hypothesis. After accounting for prior performance, behavioral engagement, study time, and number of attempts, adaptive-policy assignment retained an estimated coefficient of approximately **+3.90 achievement points**.

The model explained approximately 91.7% of the simulated outcome variance, which is expected to be high because prior achievement was deliberately incorporated as a major component of the synthetic achievement-generating mechanism. This value should therefore not be interpreted as an expected predictive performance level in actual universities.

The important result is instead that the adaptation coefficient remained positive after exposure controls were introduced.

5.3 Heterogeneity of Adaptive Benefits

The simulation further indicated that adaptation was not equally beneficial across initial achievement groups.

Students entering with lower prior achievement demonstrated an approximate adaptive uplift of **5.14 points**, compared with approximately **4.09 points** among intermediate-achievement students and **2.81 points** among initially higher-achieving students.

This gradient suggests a plausible mechanism. Students beginning with lower mastery have greater potential to benefit from appropriately targeted remediation because fixed learning sequences can expose them to material for which prerequisite knowledge is incomplete. Higher-performing students still benefit from avoiding unnecessary repetition, but ceiling effects limit the magnitude of observable score improvement.

Consequently, reporting only a single average adaptive effect may conceal meaningful variation among learners.

5.4 Educational Interpretation

The findings also highlight why high predictive accuracy should not become the sole goal of AI-supported education. The learner-state MAE was relevant because it enabled better instructional selection, but prediction was an intermediate function rather than the final objective.

An educationally valuable AI architecture should form a chain:

learner evidence → mastery estimate → instructional decision → learner response → updated learner state.

This closed-loop framework differs fundamentally from predictive analytics systems that classify students as likely to pass or fail without specifying an actionable instructional response.

The results therefore favor an interpretation of adaptive learning as a **sequential decision-support system** rather than simply a prediction system.

Limitations

Several limitations require explicit acknowledgement.

First, the reported outcome experiment is computational. Although the data structure was informed by a genuine higher-education learning analytics dataset, the learner trajectories and intervention outcomes were synthetically generated. The effect estimates must therefore not be cited as observed evidence that adaptive learning raises university grades by a specific numerical amount.

Second, the mastery state was simplified to a continuous variable. Real learning is multidimensional. A student may demonstrate strong procedural competence while retaining conceptual misconceptions within the same subject.

Third, only three instructional actions were considered. Real adaptive environments may choose among examples, videos, simulations, hints, peer activities, readings, assessment questions, or instructor interventions.

Fourth, behavioral engagement indicators cannot reveal motivation or cognitive processing directly. Long study duration can represent productive effort, confusion, distraction, or inefficient navigation.

Fifth, adaptation thresholds were deliberately interpretable and fixed. Different disciplines may require different progression criteria.

Finally, the experiment did not model possible negative behavioral consequences such as frustration with repeated remediation, excessive reliance on automated recommendations, reduced learner autonomy, or strategic manipulation of adaptive systems.

Future Scope

Future research should validate ALSIM through a preregistered randomized controlled trial in an actual higher-education course. An especially useful design would include three conditions: conventional instruction, exposure-matched digital instruction, and fully adaptive instruction. Such a design would allow the adaptation effect to be estimated independently from the general effect of increased digital practice.

A second direction involves **knowledge tracing**. Rather than representing mastery as a single value, future architectures could estimate concept-specific mastery states and their transitions over time.

Contextual bandits or reinforcement-learning methods could subsequently be used to optimize sequential instructional decisions. However, such techniques should operate within pedagogically defined constraints so that experimentation by the algorithm cannot expose students to systematically inferior learning pathways.

Future models should also integrate uncertainty. If a learner-state estimate is unreliable, the system could request a short diagnostic assessment rather than confidently making an inappropriate recommendation.

Explainable AI represents another important extension. Student-facing systems should be capable of communicating why material has been recommended—for example, because prerequisite performance indicates an unresolved concept rather than merely displaying an unexplained algorithmic decision.

Finally, fairness evaluation should examine whether adaptation quality differs systematically according to demographic background, disability status, prior educational

preparation, socioeconomic context, or technological access. Adaptive systems should reduce learning barriers rather than encode existing educational inequalities.

Conclusion

This study developed an alternative framework for examining the relationship between adaptive learning technologies and student achievement in higher education. Rather than comparing technology users with non-users, it focused on the more specific question of whether **learner-content matching contributes to achievement independently of learning exposure**.

The proposed ALSIM framework combines learner-state estimation with interpretable instructional matching across remediation, consolidation, and advancement states. A controlled fixed-seed simulation demonstrated the analytical value of separating adaptive decision quality from study duration and activity volume.

Within the computational experiment, adaptive sequencing increased instructional matching from approximately 59% to 80%, while study exposure and numbers of attempts remained essentially equivalent. The adaptive condition consequently achieved a simulated 4.20-point advantage in final performance and a higher normalized learning gain. Regression adjustment retained an estimated adaptive contribution of approximately 3.90 points after accounting for prior achievement and behavioral variables.

The results should not be interpreted as empirical effect estimates from real university students. Their contribution is methodological: they demonstrate how future adaptive-learning studies can distinguish the value of **personalization itself** from the value of additional practice.

The broader implication is that the effectiveness of AI-supported adaptive learning should not be judged according to algorithmic complexity, recommendation frequency, or platform engagement. Its educational value depends on whether the system can estimate learner needs with sufficient accuracy, select pedagogically appropriate actions, preserve learner and instructor agency, and produce measurable improvement beyond what equivalent non-adaptive instruction would have achieved.

References

1. Banihashem, S. K., Noroozi, O., van Ginkel, S., Macfadyen, L. P., & Biemans, H. J. A. (2022). A systematic review of the role of learning analytics in enhancing feedback practices in higher education. *Educational Research Review*, 37, 100489. <https://doi.org/10.1016/j.edurev.2022.100489>
2. 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>
3. Garrison, D. R. (1997). Self-directed learning: Toward a comprehensive model. *Adult Education Quarterly*, 48(1), 18–33. <https://doi.org/10.1177/074171369704800103>
4. Kuzilek, J., Hlosta, M., & Zdrahal, Z. (2017). Open University Learning Analytics dataset. *Scientific Data*, 4, 170171. <https://doi.org/10.1038/sdata.2017.171>
5. Siemens, G. (2013). Learning analytics: The emergence of a discipline. *American Behavioral Scientist*, 57(10), 1380–1400. <https://doi.org/10.1177/0002764213498851>
6. Zhong, L. (2023). A systematic review of personalized learning in higher education: Learning content structure, learning materials sequence, and learning readiness support. *Interactive Learning Environments*, 31(10), 7053–7073. <https://doi.org/10.1080/10494820.2022.2061006>
7. UNESCO. (2023). *Guidance for generative AI in education and research*. UNESCO.
8. Baker, R. S., & Inventado, P. S. (2014). Educational data mining and learning analytics. In J. A. Larusson & B. White (Eds.), *Learning analytics: From research to practice* (pp. 61–75). Springer.
9. 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.
10. VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4), 197–221.