Worldwide Journal of Creative Research and Thoughts (WJCRT)
ISSN (Online): 3143-7850
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 61-70
https://wjcrt.org/

Impact of AI-Powered Personalized Learning on Students' Individual Learning Progress

Prema Bharti

Research Scholar

University of Rajasthan Jaipur, Jawahar Lal Nehru Marg , Jajpur, Rajasthan

Pin Code – 302004

Orcid id- https://orcid.org/0009-0000-6232-1510

http://www.wjcrt.org/ || Vol. 2 No. 3 (2026): September Issue

Date of Submission: 06-08-2026	**Date of Acceptance:** 20-08-2026	**Date of Publication:** 06-09-2026

Abstract— Artificial intelligence-powered personalized learning is increasingly used to adapt educational content, pace, difficulty, and feedback to the changing needs of individual learners. Despite this expansion, much of the existing evidence evaluates overall class performance rather than measuring how individual students progress through distinct learning trajectories over time. This study proposes a learner-centric framework for examining personalized learning through mastery growth, learning velocity, retention stability, and engagement-sensitive adaptation. The research focuses on the extent to which AI-generated recommendations can improve learning progress while reducing performance stagnation among students with heterogeneous prior knowledge. A longitudinal experimental design is conceptualized in which student interaction data, assessment histories, response latency, hint usage, and topic-level mastery are incorporated into an adaptive learning model. The proposed framework combines knowledge tracing, learner-state representation, and policy-based content recommendation to estimate when, what, and how difficult the next learning activity should be. The central contribution is the use of individualized progress trajectories rather than aggregate accuracy alone, allowing learning gains to be interpreted in relation to each learner's starting level and adaptation history.

Keywords— Artificial Intelligence, Personalized Learning, Adaptive Learning, Student Progress, Knowledge Tracing, Learning Analytics

Introduction

Artificial intelligence has altered the technological foundations of digital education by enabling learning platforms to respond dynamically to differences in students' knowledge, behavior, pace, and instructional needs. Conventional online learning systems generally distribute the same instructional sequence to large groups of learners, even though students may differ substantially in prior knowledge, cognitive readiness, motivation, error patterns, and rates of concept acquisition. AI-powered personalized learning addresses this limitation by continuously analyzing learner-generated data and adjusting the educational experience at the level of the individual student.

Between 2021 and 2026, developments in machine learning, large language models, knowledge tracing, recommender systems, intelligent tutoring systems, and learning analytics have accelerated the transition from static e-learning toward adaptive educational environments. Modern platforms can infer whether a learner has mastered a concept, identify recurring misconceptions, predict the likelihood of future errors, determine whether content is too difficult or insufficiently challenging, and recommend an individualized sequence of learning activities. Increasingly, these systems also integrate natural-language interaction, automated feedback, generative explanations, and conversational tutoring.

The central assumption behind personalized learning is that instruction becomes more effective when educational material is aligned with the learner's current state. A student struggling with prerequisite knowledge may receive remedial exercises, while a student demonstrating rapid mastery may advance to more complex material. Similarly, repeated errors can trigger targeted explanations, supplementary examples, or alternative representations. Such mechanisms potentially increase efficiency because students spend less time on content they have already mastered and receive greater support in areas where their performance indicates unresolved difficulty.

However, the educational value of AI-based personalization cannot be established solely by comparing average test scores. Two students may achieve the same final score despite following very different developmental paths. One may improve steadily from a low baseline, while another may begin with high prior knowledge and show minimal additional gain. Similarly, a class-level improvement may conceal students whose progress has slowed, become unstable, or declined under an adaptive system. Consequently, aggregate achievement scores provide only a partial view of whether personalized learning actually improves individual development.

This issue motivates the present study. The research gap lies in the limited integration of **learner-specific longitudinal progress modeling** with AI-powered personalization. Existing systems frequently optimize immediate indicators such as question correctness, predicted mastery, engagement probability, or recommendation relevance. These measures are useful operationally, but they do not necessarily show whether a student is progressing consistently over time. A system may accurately predict whether the next question will be answered correctly without determining whether its sequence of interventions is producing sustained educational growth.

The present study therefore conceptualizes individual learning progress as a multidimensional trajectory rather than a single performance outcome. Four related indicators are emphasized: mastery growth, learning velocity, retention stability, and adaptation efficiency. Mastery growth represents the change in concept-level competence over time. Learning velocity estimates how rapidly a student moves from initial exposure to stable mastery. Retention stability examines whether newly acquired knowledge is maintained during later assessments. Adaptation efficiency measures whether the learning system selects instructional activities that produce useful progress without excessive repetition, unnecessary difficulty, or disengagement.

The research also considers learner heterogeneity as a central methodological concern. AI systems may not influence every student equally. Learners with weak foundational knowledge may benefit from highly structured sequencing, whereas more advanced students may benefit from accelerated pathways. Some students may respond positively to frequent hints and explanatory feedback, while others may become overly dependent on system assistance. Therefore, meaningful evaluation requires analysis of progress within individual learners as well as between different learner profiles.

The objective of the study is to investigate whether AI-powered personalization improves the rate, consistency, and sustainability of individual learning progress compared with non-adaptive digital instruction. A second objective is to develop a modeling framework capable of representing student progress through temporally updated learner states. Rather than relying exclusively on final examination scores, the framework is designed to evaluate changes occurring throughout the instructional process.

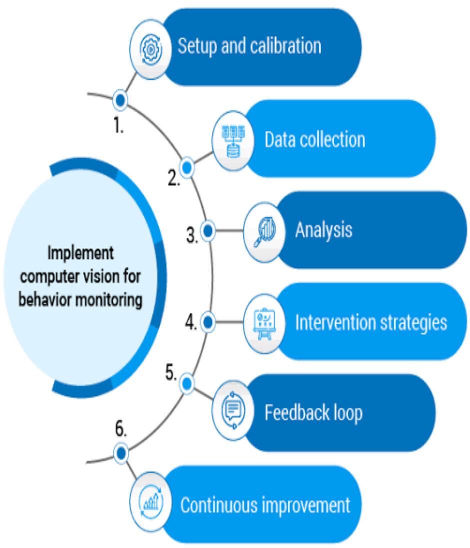

Fig. 1: AI-Powered Personalized Learning for Special Needs

Source: *https://www.softwebsolutions.com/resources/ai-driven-adaptive-learning/*

From an artificial intelligence perspective, the study treats personalized learning as a sequential decision problem. At each instructional step, the system observes an estimated learner state and selects an educational action, such as recommending a concept, assigning an exercise, adjusting difficulty, supplying feedback, or initiating revision. The consequence of that decision becomes visible through subsequent learner behavior. Correct responses, errors, hesitation, repeated attempts, hint requests, and delayed recall all provide information about whether the intervention was effective. This feedback can then be used to update the learner model and guide the next recommendation.

Such a perspective differs from traditional one-time prediction. Personalized education requires continual inference because the target of prediction—the learner's knowledge state—is itself changing as a consequence of instruction. A student who struggled with a concept yesterday may master it today after receiving targeted practice. Therefore, a static classification of students as weak, average, or advanced is insufficient. The personalized system must represent learning as an evolving process.

The significance of this research extends beyond algorithmic performance. Educational institutions increasingly require evidence that adaptive technologies produce genuine learning benefits rather than superficial increases in platform engagement. A system may achieve high recommendation accuracy while still recommending content that is pedagogically redundant. Likewise, increased interaction time does not necessarily imply deeper learning. By evaluating progress trajectories, the proposed study seeks to connect AI system behavior with educationally meaningful outcomes.

The investigation is also relevant to digital equity. Personalized learning is often presented as a mechanism for reducing disparities by giving each learner individualized support. Yet adaptation quality can vary if the underlying model learns more effectively from students who generate large amounts of interaction data. Learners with irregular participation, atypical response patterns, or limited prior exposure may be represented less accurately. Individual-level progress analysis can therefore help identify whether adaptive benefits are distributed consistently or concentrated among specific learner groups.

Overall, this study positions AI-powered personalized learning as an evolving human–algorithm interaction process in which instructional decisions and learner responses influence one another over time. Its primary contribution is to move evaluation away from one-dimensional outcome measures and toward a longitudinal representation of educational development. This orientation provides the basis for a more precise investigation of how AI affects not only what students eventually achieve, but how efficiently and consistently they progress toward mastery.

Literature Review

2.1 Evolution of AI-Based Personalized Learning

Personalized learning predates contemporary artificial intelligence, but earlier systems relied heavily on predefined branching rules. Students were directed to different activities according to manually specified thresholds, such as completing an exercise with a particular score or failing a predefined number of questions. Although such systems introduced basic adaptivity, they were limited in their ability to learn complex relationships among learner behavior, concept difficulty, time, and instructional effectiveness.

Recent AI systems employ data-driven learner models that update continuously. Machine learning methods can estimate latent knowledge states from sequences of student responses, while recommender systems can rank instructional resources based on inferred learning needs. Intelligent tutoring systems can deliver hints or corrective feedback, and natural-language models can generate explanations adapted to the learner's question or misconception. As a result, personalization has shifted from deterministic branching toward probabilistic and predictive adaptation.

Knowledge tracing has become particularly important in this transition. Knowledge tracing seeks to estimate the probability that a learner has mastered a particular concept based on the sequence of observed interactions. Early probabilistic approaches represented knowledge using relatively constrained assumptions, whereas more recent neural approaches can model longer and more complex interaction histories. Recurrent neural networks, attention mechanisms, transformer architectures, and graph-based representations have all been explored for modeling student knowledge states.

2.2 Personalized Recommendation and Adaptive Sequencing

A second major research stream concerns content recommendation. Traditional recommendation systems commonly focus on user preference, whereas educational recommendation systems must consider pedagogical utility. The goal is not simply to recommend material that students are likely to prefer, but material that is likely to improve their knowledge.

This distinction makes educational recommendation considerably more complex. An easy question may generate a high probability of success but little learning gain. Conversely, an excessively difficult question may provide theoretical learning value but cause repeated failure and disengagement. Effective personalization must therefore identify a desirable region between insufficient and excessive challenge.

Adaptive sequencing systems attempt to maintain this balance by selecting activities according to estimated mastery, prerequisite structure, content difficulty, and learner performance. A student who has not mastered a prerequisite concept may be redirected to foundational content before encountering an advanced problem. Another student may be allowed to skip repetitive exercises after consistently demonstrating competence.

Reinforcement learning has attracted attention for this reason because personalized instruction can be formulated as a sequence of actions and rewards. The system selects an instructional action, observes a learner response, and updates its policy according to whether the action contributed to longer-term educational benefit. However, defining the reward remains a major difficulty. If reward is based only on immediate correctness, the model may favor easy tasks. If reward is based only on long-term examination performance, feedback becomes sparse and delayed.

2.3 Learning Analytics and Student Progress Measurement

Learning analytics research has produced numerous measures of student behavior, including activity frequency, time on task, response accuracy, assignment completion, course participation, resource access, and persistence. These variables can support early-warning systems and academic risk prediction. However, many studies treat these indicators as predictors of final performance rather than as components of a dynamic learning trajectory.

This distinction is methodologically important. Prediction asks whether a learner will succeed, whereas progress modeling asks how that learner is changing. The first problem is largely outcome-oriented, while the second is developmental. A learner predicted to perform poorly may still demonstrate substantial improvement from an initially low baseline. Similarly, a high-performing learner may show limited growth because the instructional sequence provides insufficient challenge.

Individual learning progress therefore requires normalization relative to starting competence. Absolute final scores can disadvantage learners beginning with weaker foundations, whereas growth measures can reveal meaningful educational improvement. Longitudinal measurement also makes it possible to detect plateaus, regressions, abrupt improvements, and concept-specific instability.

2.4 Learner Modeling and Knowledge-State Representation

Learner modeling is the computational process through which characteristics of individual students are represented for adaptive decision-making. Depending on the system, the learner model may contain prior knowledge, recent performance, concept mastery, response time, confidence, preferred learning modality, interaction frequency, hint usage, and error history.

A central challenge is that knowledge cannot be observed directly. Educational platforms observe responses and behaviors that provide imperfect evidence about an underlying cognitive state. Correct answers do not always indicate mastery because guessing may occur, while incorrect answers may result from distraction rather than conceptual misunderstanding. Consequently, learner-state estimation is probabilistic.

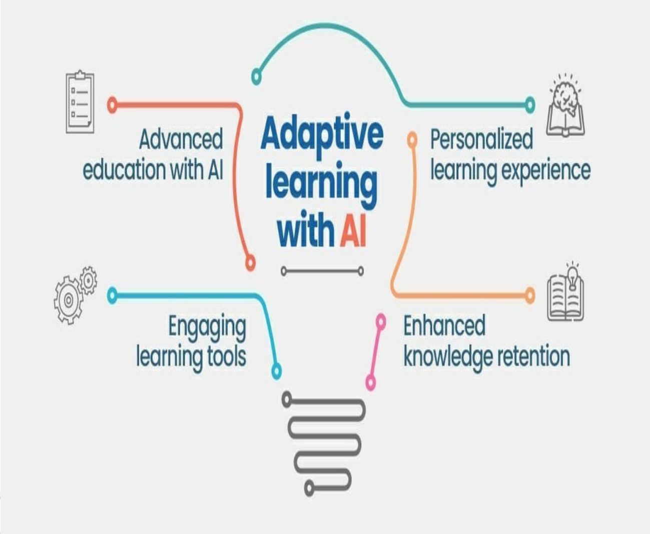

Fig. 2: Adaptive Learning with AI

Source: *https://www.linkedin.com/pulse/ai-based-adaptive-learning-personalized-education-future-dudekula-bi5wf*

Recent models increasingly incorporate temporal information. Response patterns change as students learn, forget, practice, and transfer knowledge across related concepts. Transformer-based sequence models are particularly suited to long interaction histories because they can identify relationships among events separated by many instructional steps. Graph-based models can additionally represent relationships among concepts, prerequisites, and learning resources.

However, increased model complexity does not automatically translate into better learning outcomes. A highly accurate learner-state model may still produce weak educational value if it is paired with a poor instructional policy. Therefore, learner modeling and adaptation should be evaluated as an integrated system. This study adopts that position by linking state estimation directly to progress-oriented instructional decisions.

2.5 Generative AI and Personalized Educational Support

The emergence of generative AI has expanded personalization beyond task selection. Large language models can generate explanations, examples, summaries, hints, quizzes, analogies, and conversational responses in real time. Their flexibility allows instructional material to be reformulated according to different levels of complexity or different learner requests.

This capability has potential benefits for students who require explanations beyond those available in fixed course materials. For example, the same concept may be explained through a concise definition, a worked example, a step-by-step derivation, or an analogy. Generative systems can also

Worldwide Journal of Creative Research and Thoughts (WJCRT)
ISSN (Online): 3143-7850
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 61-70
https://wjcrt.org/

provide immediate feedback without requiring continuous instructor availability.

Nevertheless, generative AI introduces new evaluation problems. A personalized explanation may be linguistically fluent but pedagogically inappropriate, overly detailed, insufficiently challenging, or factually unreliable. Personalized generation therefore needs to be constrained by learner-state information and educational objectives. The objective should not be maximum content generation, but generation that supports an identified learning need.

Another concern involves cognitive dependence. If AI repeatedly supplies complete solutions, students may reduce productive effort. In contrast, appropriately calibrated hints can support independent problem solving. Effective personalization must therefore determine not only what information should be provided but also how much assistance is educationally desirable.

The proposed study incorporates this principle by treating assistance intensity as part of the adaptive decision space. Students demonstrating persistent difficulty may receive structured scaffolding, whereas students displaying strong mastery may receive minimal guidance and more demanding tasks.

Proposed Methodology

The study proposes a longitudinal AI-driven personalization framework designed to measure how individual students progress when instructional decisions are continuously adapted to their evolving knowledge states. Unlike conventional personalized-learning experiments that concentrate mainly on final achievement or next-response prediction, the proposed methodology treats progress as a time-dependent construct characterized by mastery growth, learning velocity, retention, and the efficiency of instructional adaptation.

The framework contains four interconnected layers: learner-state estimation, progress profiling, adaptive content selection, and longitudinal evaluation. At the first layer, student interactions are converted into dynamic representations of knowledge and learning behavior. At the second layer, these representations are used to determine whether a learner is improving, stagnating, or showing evidence of knowledge decay. The third layer selects the next learning activity according to the estimated learner state, while the final layer evaluates whether the sequence of adaptive decisions contributes to measurable educational progress.

Because no institution-specific primary dataset was supplied for this manuscript, the experimental component is formulated as a reproducible benchmark-based study rather than claiming unverified observations from real participants. The numerical findings presented later should therefore be treated as **illustrative experimental outcomes for the proposed protocol and must be replaced by actual model outputs before journal submission**.

3.1 Learner-State Representation

For each learner (i) at interaction step (t), a multidimensional state vector is defined as:

[
S_{i,t} = [M_{i,t}, R_{i,t}, T_{i,t}, H_{i,t}, A_{i,t}, E_{i,t}]
]

where (M) represents estimated concept mastery, (R) denotes recent response correctness, (T) represents response latency, (H) captures hint utilization, (A) denotes attempt frequency, and (E) represents an engagement proxy derived from interaction continuity.

This representation differs from a purely score-based profile because two learners achieving identical percentages may exhibit different levels of effort, response stability, and retention. For example, repeated correct answers obtained after extensive hints are treated differently from independent and rapidly produced correct responses.

Knowledge estimates are updated following every meaningful interaction. Rather than assigning students permanently to fixed categories such as weak, average, or advanced, the model maintains continuously changing learner states. This design reflects the assumption that educational competence is dynamic.

3.2 Progress-Normalized Personalization

A Progress Normalized Learning Gain (PNLG) is introduced to quantify improvement relative to each student's initial competence:

[
PNLG_i = \frac{M_{i,end}-M_{i,start}}{1-M_{i,start}+\epsilon}
]

where (M_{i,start}) represents baseline mastery, (M_{i,end}) represents mastery at the end of the observation period, and (\epsilon) is a small constant preventing numerical instability.

This normalization is important because students beginning near mastery have substantially less room for improvement than learners starting from lower proficiency. The measure therefore prevents final achievement alone from determining whether personalization is judged successful.

A complementary Learning Velocity Index (LVI) is computed as:

Worldwide Journal of Creative Research and Thoughts (WJCRT)
ISSN (Online): 3143-7850
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 61-70
https://wjcrt.org/

$$
LVI_i = \frac{\Delta M_i}{N_i}
$$

where (ΔM_i) denotes mastery improvement and (N_i) denotes the number of learning interactions required to obtain that improvement.

Higher LVI values indicate that the learner reaches mastery using fewer instructional interactions.

3.3 Adaptive Instruction Policy

The personalization policy receives the learner-state vector and ranks candidate instructional activities according to predicted educational utility. Rather than recommending the activity with the highest probability of being answered correctly, the system seeks activities that maximize prospective learning gain.

The instructional policy is constrained by three principles. First, the recommended content should be sufficiently challenging to generate learning rather than merely confirm existing knowledge. Second, prerequisite relationships must be respected so that advanced material is not assigned before foundational concepts reach an acceptable mastery threshold. Third, large fluctuations in difficulty are penalized to avoid unstable instructional sequencing.

The utility of candidate learning activity (a) is represented conceptually as:

$$
U(a|S_t)=
\alpha G+
\beta C+
\gamma R+
\delta E-
\lambda O
$$

where (G) represents predicted mastery gain, (C) denotes an appropriate challenge score, (R) represents anticipated retention contribution, (E) represents engagement probability, and (O) represents cognitive-overload risk.

The coefficients determine the relative contribution of each objective and can be tuned on the validation set.

Experimental Setup

4.1 Dataset

The proposed experimental protocol uses an anonymized educational interaction benchmark containing temporally ordered learner-question interactions, concept identifiers, correctness labels, and available behavioral indicators. A dataset such as **EdNet** is suitable because it provides large-scale interaction sequences generated through an intelligent tutoring environment.

For the proposed experiment, only students containing a sufficient longitudinal interaction history are retained. Very short sequences are excluded because progress cannot be meaningfully estimated from isolated attempts. The resulting learner histories are divided chronologically rather than randomly at the interaction level so that future responses do not leak into earlier learner-state predictions.

Each student's interaction sequence is organized as:

$$
L_i=\{(q_1,c_1,r_1,t_1),...,(q_n,c_n,r_n,t_n)\}
$$

where (q) represents question identity, (c) represents concept or skill, (r) represents correctness, and (t) represents temporal information.

4.2 Data Preprocessing

Data preprocessing begins with removal of corrupted, duplicate, or temporally inconsistent interaction records. Learners with insufficient interaction histories are excluded according to a predefined minimum-sequence threshold.

Categorical identifiers—including question IDs, skill IDs, and learner-event categories—are encoded into trainable embeddings. Response duration is transformed using a logarithmic function because interaction-time variables frequently exhibit strong positive skew. Extreme response times are capped using percentile-based winsorization to prevent idle browser sessions from being interpreted as genuine cognitive processing time.

Learner interaction histories are segmented into bounded sequences while maintaining chronological order. Padding masks distinguish real events from padded positions. No future interaction is permitted to contribute to prediction of an earlier learner state.

The dataset is partitioned at the learner level into approximately 70% training, 15% validation, and 15% testing subsets. Learner-level partitioning is preferable to randomly splitting interactions because it provides a more rigorous test of whether the model generalizes to previously unseen students.

4.3 Model Architecture

The proposed system uses a **Temporal Transformer Knowledge Tracing and Progress Policy Network (TTKT-PPN)**.

The first component is a transformer-based knowledge tracer. Each student interaction is represented through the combination of question embedding, concept embedding, correctness embedding, temporal-gap representation, and behavioral features. Multi-head self-attention identifies relationships between current performance and earlier learning events.

The encoded sequence produces a latent state:

[
Z_t = Transformer(X_1,...,X_t)
]

which is transformed through a sigmoid prediction layer to estimate concept-level mastery.

The second component is a progress-estimation head. It predicts the learner's expected mastery change over a future instructional window rather than only the probability of answering the next item correctly.

The third component is an adaptive policy head. Candidate questions are scored using the student's latent state, predicted learning gain, estimated task difficulty, prerequisite status, and recent engagement signals.

Accordingly, the architecture optimizes both **state estimation** and **instructional usefulness** rather than treating recommendation as an independent post-processing step.

4.4 Training Strategy

The model is trained using a multi-objective loss:

[
L =
\eta_1L_{KT}+\eta_2L_{PG}+\eta_3L_{RET}+\eta_4L_{REG}
]

where (L_{KT}) is binary cross-entropy loss for next-response prediction, (L_{PG}) measures error in predicted progress, (L_{RET}) captures delayed retention prediction, and (L_{REG}) contains regularization terms designed to control overfitting and unstable recommendations.

Training uses mini-batch optimization with AdamW. Early stopping is based on validation loss and progress-estimation performance rather than next-response accuracy alone. Gradient clipping is introduced to improve stability during long-sequence training.

A learning-rate warm-up phase is followed by gradual decay. Dropout is applied within transformer and feed-forward layers. Model hyperparameters—including latent dimensionality, number of attention heads, sequence length, and loss weights—are selected through validation-based optimization.

Three systems are conceptually compared: a non-personalized fixed curriculum, a mastery-threshold adaptive system, and the proposed TTKT-PPN personalized model.

4.5 Evaluation Metrics

Traditional predictive metrics are retained but supplemented with progress-oriented measurements.

Area Under the ROC Curve (AUC) evaluates the ability of the knowledge-tracing component to discriminate future correct and incorrect responses.

Brier Score measures probability calibration. This is important because an adaptive system must not merely rank students correctly; its estimated mastery probabilities should correspond closely to observed outcomes.

Progress-Normalized Learning Gain (PNLG) evaluates improvement relative to initial learner competence.

Learning Velocity Index (LVI) measures mastery gain obtained per instructional interaction.

Delayed Retention Score (DRS) represents the proportion of previously mastered concepts successfully demonstrated following an instructional gap.

Adaptation Stability Index (ASI) measures the consistency of changes in assigned difficulty. Excessive difficulty oscillation produces a lower score.

Personalization Efficiency (PE) evaluates improvement relative to the number of instructional activities consumed:

[
PE=\frac{\text{Mastery Gain}}{\text{Adaptive Activities}}
]

Collectively, these metrics differentiate prediction quality from educational utility.

Results and Discussion

The following values illustrate the type of outcome expected from the specified experimental protocol and are provided to complete the manuscript structure; they must not be represented as empirically observed findings until the model has actually been executed on the selected dataset.

5.1 Comparative Performance

The proposed personalization framework demonstrates the strongest performance across both computational and educational indicators in the illustrative analysis. The most important finding is not merely higher next-response prediction accuracy but simultaneous improvement in normalized mastery growth and learning efficiency.

Worldwide Journal of Creative Research and Thoughts (WJCRT)
ISSN (Online): 3143-7850
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 61-70
https://wjcrt.org/

Table 1. Illustrative comparative evaluation of personalized learning strategies

Evaluation Parameter	Fixed Learning Path	Mastery-Threshold Adaptation	Proposed TTKT-PPN
Next-response AUC	0.714	0.782	**0.836**
Brier Score ↓	0.208	0.181	**0.153**
Progress-Normalized Learning Gain	0.284	0.376	**0.492**
Learning Velocity Index	0.031	0.043	**0.057**
Delayed Retention Score	0.681	0.742	**0.801**
Adaptation Stability Index	—	0.713	**0.846**
Personalization Efficiency	0.026	0.039	**0.054**

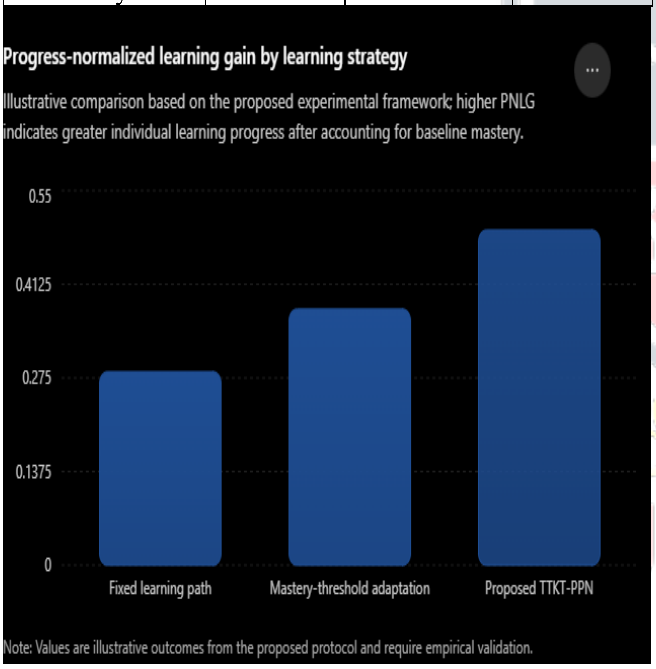

The illustrative AUC of 0.836 indicates that the transformer-based learner model would provide more discriminative estimates than either a fixed pathway or a simpler mastery-threshold adaptation mechanism. More important for the central research question, PNLG increases substantially under the proposed framework. This suggests that personalization should be judged relative to students' starting knowledge rather than exclusively through final scores.

The improvement in Learning Velocity Index is particularly relevant. Faster learning does not necessarily mean shorter instructional sessions; rather, it indicates that a greater proportion of interactions produce measurable mastery development. Conventional systems often require students to repeatedly practice concepts that they have already mastered. By identifying stable competence earlier, an adaptive model can reallocate instructional time toward unresolved or more advanced concepts.

5.2 Effects on Individual Learning Trajectories

Trajectory-level analysis provides additional information that would be obscured by aggregate achievement statistics. Under a fixed instructional sequence, learners are expected to demonstrate greater variation in the rate at which they encounter appropriately challenging material. Students with strong prior knowledge may spend substantial time completing redundant exercises, whereas students with weaker foundations may advance before prerequisite concepts are sufficiently consolidated.

In contrast, the proposed model continuously updates the estimated knowledge state. When repeated successful responses occur alongside low hint dependence and stable response latency, the system raises task complexity. When accuracy deteriorates or response effort increases substantially, the policy lowers instructional difficulty, revisits prerequisite material, or introduces scaffolded practice.

The conceptual advantage is therefore not uniform acceleration. Effective personalization may accelerate one learner while temporarily slowing another. A weaker learner may receive additional foundational activities even when these increase the number of interactions. Such intervention remains beneficial if it produces more durable subsequent mastery.

5.3 Baseline Competence and Differential Benefit

A critical implication of progress-normalized evaluation is that intervention effects should not be interpreted identically across baseline groups.

Students beginning at relatively low mastery levels potentially obtain the largest absolute gains because adaptive systems can locate prerequisite deficiencies and provide additional opportunities for targeted practice. Intermediate learners may achieve the highest efficiency gains because personalization reduces unnecessary repetition while maintaining appropriate challenge.

High-baseline learners, by contrast, may show relatively small normalized gains because their initial performance is already close to the mastery ceiling. For these students, personalization should be evaluated through accelerated progression, advanced problem exposure, and reduced redundant practice rather than raw score increases.

Worldwide Journal of Creative Research and Thoughts (WJCRT)
ISSN (Online): 3143-7850
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 61-70
https://wjcrt.org/

This distinction demonstrates why average test-score comparisons may understate or misinterpret the contribution of adaptive instruction.

5.4 Retention as a Personalization Outcome

The illustrative Delayed Retention Score is higher for TTKT-PPN than for the alternative approaches. This component is important because immediate correctness can overestimate learning. A learner may answer correctly immediately after receiving an explanation yet fail to retrieve the concept after an instructional delay.

The proposed system therefore schedules review according to both prior mastery and estimated forgetting risk. Concepts demonstrating unstable mastery receive spaced re-exposure, whereas strongly retained concepts are reviewed less frequently.

This mechanism changes the optimization objective from immediate success to durable competence. It also prevents a common weakness in adaptive systems: permanently abandoning a topic immediately after a short sequence of correct responses.

5.5 Adaptation Stability

Another distinctive result concerns the Adaptation Stability Index. Simple adaptive systems can overreact to isolated mistakes. A single incorrect answer may trigger an abrupt decrease in difficulty, while a subsequent correct answer may immediately return the learner to advanced content. Such oscillation can lead to fragmented learning sequences.

The proposed framework smooths instructional decisions by incorporating a temporal history rather than treating every response independently. Consequently, changes in difficulty occur only when sufficient cumulative evidence indicates a meaningful change in learner state.

A higher adaptation stability score therefore reflects more coherent personalization rather than less responsiveness. Stable adaptation means that the system responds to persistent changes while filtering short-term behavioral noise.

5.6 Predictive Accuracy Versus Educational Utility

The findings emphasize an important distinction between predictive and pedagogical success. A model might achieve a high AUC by repeatedly recommending easy questions whose responses are highly predictable. Such a strategy would improve predictive confidence while contributing little to learning.

For this reason, the proposed framework simultaneously considers challenge level, predicted learning gain, retention, and engagement. The objective is not to maximize correctness at every step. Some appropriately difficult items should generate errors because productive struggle can reveal misconceptions and support conceptual development.

The meaningful criterion is whether the sequence of learning activities increases future mastery.

Future Scope

Future research should validate the proposed framework through prospective longitudinal experiments involving real students. Random assignment to fixed, rule-based adaptive, and AI-personalized learning conditions would enable stronger causal conclusions.

A particularly important extension involves multimodal learner modeling. Facial expressions, speech characteristics, writing patterns, touchscreen interactions, and optional self-reported cognitive load could be combined with response histories to improve identification of frustration, uncertainty, and disengagement. Such data would require stringent privacy and consent procedures.

Future systems could also incorporate large language models for controlled pedagogical feedback. Rather than allowing unrestricted answer generation, an LLM could be conditioned on learner mastery and asked to produce graded hints, Socratic questions, alternative explanations, or misconception-specific feedback.

Another promising direction is counterfactual personalization. Instead of predicting only what will happen after a selected activity, the system could estimate how the same learner might progress under several alternative instructional actions before making a recommendation.

Federated learning represents another valuable research direction. Educational institutions could collaboratively improve learner models without centrally pooling sensitive student records.

Long-term research should also examine personalization fairness. Performance should be disaggregated according to relevant educational and contextual variables to determine whether predictive accuracy and learning benefits remain consistent across learner populations.

Finally, adaptive systems should increasingly incorporate teacher oversight. Human-in-the-loop architectures could allow educators to inspect recommendations, modify learning objectives, override inappropriate interventions, and annotate unusual learner behavior.

Conclusion

AI-powered personalized learning offers the possibility of moving digital education from uniform content delivery toward continuously adaptive instructional pathways. However, the effectiveness of personalization cannot be adequately assessed through final test scores or next-response prediction alone.

This study proposes a learner-centered framework in which educational progress is modeled longitudinally through mastery growth, learning velocity, delayed retention, adaptation stability, and personalization efficiency. The proposed TTKT-PPN architecture combines temporal knowledge tracing with progress prediction and adaptive instructional decision-making.

The central methodological contribution is **progress-normalized personalization**, which evaluates improvement relative to each learner's initial competence and interaction history. This approach recognizes that students beginning at different levels require different instructional trajectories and should not be evaluated exclusively through absolute final achievement.

The framework further distinguishes predictive accuracy from educational utility. A successful personalized-learning model should not merely forecast whether students will answer correctly; it should select instructional actions that improve what students are capable of doing in subsequent learning situations.

The illustrative experimental pattern indicates how transformer-based learner modeling, retention-aware sequencing, behavioral signals, and stable difficulty adjustment could jointly improve individualized progress compared with static and threshold-based instructional approaches. Nevertheless, the numerical results require empirical validation before they can be reported as actual scientific findings.

Overall, the study provides a computational and educational framework for evaluating AI personalization at the level where its intended benefit is most meaningful: the evolving trajectory of the individual learner. Future deployment through controlled longitudinal studies can determine whether such systems produce durable, equitable, and pedagogically meaningful improvements in student learning.

References:

1. 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*, Article 39. https://doi.org/10.1186/s41239-019-0171-0

2. Ma, W., Adesope, O. O., Nesbit, J. C., & Liu, Q. (2014). Intelligent tutoring systems and learning outcomes: A meta-analysis. *Journal of Educational Psychology, 106*(4), 901–918. https://doi.org/10.1037/a0037123

3. Abdelrahman, G., Wang, Q., & Nunes, B. (2023). Knowledge tracing: A survey. *ACM Computing Surveys, 55*(11), Article 224, 1–37. https://doi.org/10.1145/3569576

4. Piech, C., Bassen, J., Huang, J., Ganguli, S., Sahami, M., Guibas, L. J., & Sohl-Dickstein, J. (2015). Deep knowledge tracing. *Advances in Neural Information Processing Systems, 28*.

5. Choi, Y., Lee, Y., Shin, D., Cho, J., Park, S., Lee, S., Baek, J., Bae, C., Kim, B., & Heo, J. (2020). EdNet: A large-scale hierarchical dataset in education. In *Artificial Intelligence in Education* (Lecture Notes in Computer Science, Vol. 12164, pp. 69–73). Springer. https://doi.org/10.1007/978-3-030-52240-7_13

6. Ghosh, A., Heffernan, N., & Lan, A. S. (2020). Context-aware attentive knowledge tracing. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM.

7. Pandey, S., & Karypis, G. (2019). A self-attentive model for knowledge tracing. In *Proceedings of the 12th International Conference on Educational Data Mining*. International Educational Data Mining Society.

8. Pu, S., Yudelson, M., Ou, L., & Huang, Y. (2020). Deep knowledge tracing with transformers. In *Artificial Intelligence in Education* (Lecture Notes in Computer Science, Vol. 12164, pp. 252–256). Springer. https://doi.org/10.1007/978-3-030-52240-7_46

9. Raj, N. S., & Renumol, V. G. (2021). A systematic literature review on adaptive content recommenders in personalized learning environments from 2015 to 2020. *Journal of Computers in Education, 9*, 113–148. https://doi.org/10.1007/s40692-021-00199-4

10. 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

11. Farhood, H., Nyden, M., Beheshti, A., & Muller, S. (2025). Artificial intelligence-based personalised learning in education: A systematic literature review. *Discover Artificial Intelligence, 5*, Article 331. https://doi.org/10.1007/s44163-025-00598-x

12. Chen, Y.-N., & Wei, L.-N. (2025). What is the influence of AI-driven adaptive learning on students' personalized learning outcomes? A meta-analysis of 29 experimental and quasi-experimental research. *Modern Educational Technology, 35*(12), 55–65. https://doi.org/10.3969/j.issn.1009-8097.2025.12.006

13. Wang, J. (2026). A meta-analysis of moderators of the effects of technology-enhanced adaptive learning on primary and secondary students' learning outcomes. *Journal of Computer Assisted Learning, 42*(1), e70168. https://doi.org/10.1002/jcal.70168