IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 7, July 2026

DOI: 10.17148/IJARCCE.2026.15744

AI-Based Personalized Learning System for Students: An Adaptive Framework for Learner Modeling and Content Recommendation

Atul Sharma[1], Aashish Kumar Tiwari[2]

Department of Computer Science and Engineering, SAM Global University, Raisen, India[1]

Department of Computer Science and Engineering, SAM Global University, Raisen, India[2]

Abstract: Conventional classroom instruction and generic e-learning platforms tend to deliver the same content, pace, and sequence to every learner, even though students differ widely in prior knowledge, learning speed, motivation, and preferred mode of engagement. This gap between uniform delivery and diverse learner needs often results in disengagement, uneven learning outcomes, and higher dropout rates in online courses. This paper presents the design of an AI-based personalized learning system that continuously observes a student's interaction data — quiz performance, time-on-task, click patterns, and self-reported preferences — to build and update a dynamic learner profile. The profile drives a hybrid recommendation engine that combines content-based filtering with a knowledge-tracing model to select the next best learning resource, adjust difficulty, and time interventions such as hints or remedial material. A prototype covering the learner-modeling and recommendation modules was implemented using a Python-based machine learning pipeline with a web front end, and its behavior was evaluated on a simulated cohort of learners against a static, non-adaptive baseline. The adaptive pipeline produced a noticeably better fit between recommended difficulty and learner ability while reducing the number of redundant exercises presented to already-competent students. The paper also discusses the practical barriers to deploying such a system at scale, including cold-start profiling, algorithmic bias, and data privacy, and outlines directions for future work such as affect-aware adaptation and teacher-in-the-loop oversight.

Keywords: personalized learning, adaptive e-learning, artificial intelligence in education, learner modeling, knowledge tracing, recommendation systems, intelligent tutoring

I. INTRODUCTION

Classrooms have always contained a spread of abilities. A teacher standing in front of thirty students is, in effect, being asked to teach thirty different learning trajectories at once, and in practice the lecture is paced for an imagined average student who rarely exists. The shift to online and blended learning over the last decade, accelerated further during the COVID-19 disruption, moved a large share of instruction onto digital platforms, but most of these platforms inherited the same one-size-fits-all delivery model — a fixed sequence of videos, readings, and quizzes shown to every enrolled learner regardless of how they are actually doing [1], [6].

Personalized learning attempts to close this gap by tailoring what a learner sees, in what order, and at what pace, based on evidence collected about that specific learner rather than assumptions about the class as a whole. Artificial intelligence is well suited to this problem because it can process behavioral signals — response times, error patterns, navigation choices — at a scale and speed no instructor could manage manually, and can update its estimate of a learner's state after every interaction rather than only at test time [2], [3].

This paper describes the design and a working prototype of an AI-based personalized learning system aimed at school and early undergraduate students. The system is built around three ideas: (1) a learner model that is continuously updated rather than fixed at enrollment, (2) a recommendation engine that treats content selection as a sequential decision problem rather than a one-off classification, and (3) a feedback loop that surfaces the same information to the teacher, so personalization does not remove the human from the loop but augments their judgment.

The remainder of the paper is organized as follows. Section II reviews related work on adaptive and AI-driven learning systems. Section III presents the proposed system architecture. Section IV describes the implementation and the algorithms used. Section V reports results from an evaluation against a non-adaptive baseline. Section VI discusses limitations and open challenges, and Section VII concludes with directions for future work.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁙ Peer-reviewed & Refereed journal ⁙ Vol. 15, Issue 7, July 2026

DOI: 10.17148/IJARCCE.2026.15744

II. RELATED WORK

Adaptive e-learning has been studied from several angles, and it is useful to group prior work by the signal the system relies on to personalize instruction.

A. Learning-Style and Prior-Knowledge Based Adaptation

An early and still common approach classifies learners by a learning-style model — visual, auditory, read/write, or kinesthetic — and by their prior knowledge level, then routes them through content authored for that combination. The Adaptivo system, for instance, adjusts both the learning path and the presentation format according to a learner's detected style and knowledge level, and additionally factors in behavioral differences such as time spent and frequency of interaction [1]. A related platform aimed at K-12 learners, APPEAL, combines a VARK-style presentation layer with gamification and a difficulty-scaffolding mechanism that lets a learner skip, hide, or reattempt exercises depending on how a dynamic learner model rates their cognitive state [6]. These systems demonstrate that even a relatively simple, rule-based notion of learner type can produce a meaningfully different — and more engaging — experience than static content.

B. Machine-Learning-Based Learner Modeling

A second line of work replaces hand-authored rules with models trained on learner interaction data. A systematic review covering studies from 2015 to 2022 found that machine-learning techniques are increasingly used to infer a learner's style dynamically from behavioral traces rather than from a one-time questionnaire, which addresses the well-known problem that self-reported learning styles are often unstable and only weakly predictive of actual performance [3]. Other systems fold in richer signals: one recent platform pairs an AI-driven virtual instructor with webcam-based engagement monitoring so that the system can react not only to what a learner answers but to whether they appear to be paying attention at all [2]. Still other work has gone further and incorporated physiological measurements — wearable-sensor indicators of fatigue and cognitive load — into the model that decides how difficulty should evolve session to session, using both rule-based heuristics and supervised models trained on these physiological signals [4].

C. Feedback, Engagement, and Content Sequencing

Beyond modeling the learner, several studies focus on how the system should act once it has an estimate of learner state. Work on multiple-intelligence-aware feedback found that the way feedback is phrased and timed measurably shifts learner behavior, not just outcomes on the next quiz [5]. Content-sequencing has also been framed as a reinforcement-learning problem: the APPEAL platform uses a Deep Q-Network to decide, at each step, whether to present easier practice, harder practice, or a change in presentation mode, treating the whole learning session as a Markov decision process rather than a lookup table [6]. Generative approaches have also been explored for producing the learning content itself and for stitching individual resources into a coherent pathway, rather than personalizing only the order of a fixed catalog [13].

D. Broader Reviews and Open Concerns

Several recent surveys take stock of the field as a whole. A 2025 systematic review of 142 empirical studies concluded that machine learning, deep learning, and multimodal analytics are now routinely combined to adapt content in real time, but also flagged that pedagogical grounding and equity considerations often lag behind the technical sophistication of these systems [7]. A cross-country review comparing deployments in the United States, China, and India argued for keeping a human somewhere in the loop rather than fully automating instructional decisions [11]. A widely cited framing paper on AI in education lists curriculum vision, teacher role, and research infrastructure as the four areas still holding back broader adoption, echoing concerns raised elsewhere about data privacy, algorithmic bias, and the practical burden personalization places on teachers who must interpret system output [10], [12]. Reported effect sizes vary considerably across studies, with some reviews citing reductions in required learning time on the order of 30–50% alongside learning-outcome gains of roughly 15–25% relative to non-adaptive instruction, though these figures depend heavily on subject, population, and the specific adaptation technique used [9].

The system proposed in this paper draws on this literature but differs in emphasis: rather than committing to a single adaptation signal (style, physiology, or engagement alone), it combines lightweight behavioral telemetry with a knowledge-tracing model, and it is designed from the outset to expose its reasoning to a teacher dashboard so that automated recommendations remain auditable.

DOI: 10.17148/IJARCCE.2026.15744

III. PROPOSED SYSTEM ARCHITECTURE

The proposed system is organized into five layers, described below and summarized in Table I.

A. Data Acquisition Layer

This layer captures raw signals from the learning interface: answers submitted, time taken per question, number of hint requests, navigation between resources, and an optional short preference questionnaire completed at sign-up. Signals are timestamped and streamed to the profiling layer rather than batched, so the learner model can update within the same session rather than only after a course finishes.

B. Learner Profiling Module

Raw events are converted into a compact learner state: an estimated mastery probability per topic (using Bayesian Knowledge Tracing as a baseline and a recurrent-network-based Deep Knowledge Tracing model for comparison), an engagement score derived from session frequency and time-on-task, and a coarse preference vector describing format preference (video, text, interactive exercise). Unlike a static learning-style label, this profile is a probability distribution that shifts after every graded interaction.

C. Recommendation and Adaptation Engine

Given the current profile, this module chooses the next resource using a hybrid of content-based filtering — matching resource metadata (topic, difficulty, format) to the learner's profile — and collaborative signals drawn from how similar learners performed on the same resource. A rule-based fallback intervenes when the profile is sparse, which is the common case for a newly enrolled learner (the cold-start problem), before enough interaction data exists to trust the learned model.

D. Content Repository

Learning resources — video segments, reading passages, worked examples, and auto-graded exercises — are tagged with topic, estimated difficulty, and format at ingestion time, either by an instructor or by a lightweight NLP-based tagging pipeline that extracts keywords and estimates readability.

E. Feedback and Dashboard Layer

Students see a progress view showing mastery per topic and next recommended activity. Teachers see an aggregated dashboard flagging learners who are falling behind, are being routed disproportionately to remedial content, or whose engagement score has dropped, so that a human can intervene when the automated recommendation is not enough.

Layer	Primary Input	Primary Output
Data Acquisition	Clicks, answers, timing	Event stream
Learner Profiling	Event stream	Mastery + engagement vector
Recommendation Engine	Learner profile	Ranked resource list
Content Repository	Author-tagged resources	Searchable resource index
Dashboard	Profile + recommendation logs	Student/teacher views

TABLE I. SYSTEM LAYERS AND DATA FLOW

IV. IMPLEMENTATION

A. Dataset and Preprocessing

For prototyping, interaction logs were simulated for 500 synthetic learners across 40 topics in an introductory programming course, generated so that learner ability and topic difficulty followed a standard item-response-theory structure, with added noise to mimic guessing and slips. Each simulated session produced a record of the topic attempted, correctness, response time, and hint usage, which mirrors the schema an actual deployment would collect from a learning management system.

B. Modeling Pipeline

Bayesian Knowledge Tracing was implemented as the baseline mastery estimator, updating a per-topic mastery probability after every attempt using four standard parameters — prior knowledge, learning rate, guess, and slip. A Deep Knowledge Tracing variant, using a single-layer LSTM over the interaction sequence, was trained on 70% of the

IJARCCE

ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⧖ Peer-reviewed & Refereed journal ⧖ Vol. 15, Issue 7, July 2026

DOI: 10.17148/IJARCCE.2026.15744

simulated learners and validated on the remaining 30% to compare against the Bayesian baseline. The recommendation engine ranked candidate resources by combining the negative distance between resource difficulty and current mastery estimate with a small bonus for resources matching the learner's stated format preference, and clipped its recommendation to avoid repeatedly assigning material a learner had already mastered.

C. Tools and Stack

The pipeline was implemented in Python using scikit-learn for the Bayesian baseline and TensorFlow/Keras for the LSTM-based knowledge-tracing model. A Flask API exposed profile and recommendation endpoints, consumed by a lightweight React front end for the student and teacher views. Learner events and profiles were stored in MongoDB, chosen for its flexibility in handling the semi-structured event schema.

D. Evaluation Metrics

Two families of metrics were used: predictive metrics for the knowledge-tracing models (AUC and accuracy of next-attempt correctness prediction), and simulated-learning metrics for the full pipeline (number of exercises to reach a mastery threshold, proportion of recommendations rated as "too easy" or "too hard" against the ground-truth simulated ability, and a proxy engagement measure based on session continuation rate).

V. RESULTS AND DISCUSSION

Table II compares the adaptive pipeline against a static baseline that presents topics in a fixed, instructor-defined order regardless of learner performance, evaluated on the same simulated cohort.

Metric	Static Baseline	Adaptive System
Exercises to reach mastery (avg.)	14.6	9.8
Recommendations rated mismatched	38%	16%
Session continuation rate	61%	77%
Knowledge-tracing AUC (DKT)	—	0.81

TABLE II. SIMULATED-COHORT COMPARISON, STATIC VS. ADAPTIVE

The adaptive system reduced the average number of exercises required to reach the mastery threshold by roughly a third, and cut the proportion of recommendations judged mismatched to the simulated learner's true ability from 38% to 16%. The session continuation rate — a proxy for engagement, measuring whether a simulated learner "returned" for a subsequent session under a simple dropout model tied to mismatch frequency — improved from 61% to 77%. These gains are broadly consistent with the direction, though not necessarily the exact magnitude, of improvements reported in the wider literature on adaptive systems [9].

The LSTM-based Deep Knowledge Tracing model reached an AUC of 0.81 on held-out learners for predicting next-attempt correctness, a modest but meaningful improvement over the Bayesian baseline's tendency to underfit learners whose performance varied a great deal across topics. This is consistent with earlier findings that sequence models capture cross-topic transfer effects that a per-topic Bayesian model cannot [3].

It is worth being explicit that these results come from a simulated cohort rather than real students, so they should be read as evidence that the pipeline behaves sensibly and internally-consistently, not as a validated claim about real-world learning gains. A pilot with an actual class is a necessary next step before any deployment claim can be made.

VI. CHALLENGES AND FUTURE SCOPE

A number of practical issues need attention before a system of this kind can be deployed responsibly at scale.
First, cold start remains a real limitation: a newly enrolled learner has no interaction history, and early recommendations necessarily rely on coarse defaults, which risks a poor first impression precisely when engagement matters most. Second, any model trained on historical interaction data can encode and amplify existing bias — for example, systematically under-recommending advanced material to learners from groups that historically had less access to preparatory resources — which means fairness auditing has to be a standing part of the pipeline, not a one-time check [7], [10]. Third, the data these systems require — fine-grained behavioral traces, and in some proposed designs even physiological or webcam data [2], [4] — raises real privacy concerns, particularly for learners who are minors; data minimization and clear retention limits should be treated as design constraints rather than an afterthought.

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 7, July 2026
DOI: 10.17148/IJARCCE.2026.15744

Fourth, teacher acceptance is not automatic: a dashboard that surfaces model output without explaining the reasoning behind a recommendation is likely to be ignored or mistrusted, so future versions of the recommendation engine should expose a short, human-readable justification alongside each suggestion.

Looking ahead, three directions seem most promising. Affect-aware adaptation — incorporating signals such as frustration or boredom detected from interaction patterns rather than intrusive sensors — could allow the system to intervene before a learner disengages rather than after. Multilingual and low-bandwidth support would extend the approach to learners outside well-resourced institutions, addressing the equity gap noted in recent reviews [7], [11]. Finally, a structured teacher-in-the-loop mode, where the system proposes but a teacher can override and that override is fed back as a training signal, would combine the scale of automated personalization with the judgment a purely algorithmic system still lacks.

VII. CONCLUSION

This paper presented the design and a working prototype of an AI-based personalized learning system that models learners continuously rather than statically, and combines knowledge tracing with a hybrid recommendation engine to select appropriate content. On a simulated cohort, the adaptive pipeline reduced the number of exercises needed to reach mastery, lowered the rate of mismatched recommendations, and improved a proxy engagement measure relative to a static baseline. At the same time, cold-start behavior, fairness, privacy, and teacher trust remain open problems that determine whether such a system is actually useful in a real classroom rather than only in simulation. Future work will focus on piloting the system with real learners, adding affect-aware signals, and building the teacher-override mechanism described above, with the aim of keeping the human instructor firmly in the loop as personalization scales.

REFERENCES

[1]. “Adaptivo: A Personalized Adaptive E-Learning System based on Learning Styles and Prior Knowledge,” in Proc. IEEE Conf. Publication, 2022.
[2]. “Personalized Adaptive Learning Platform Empowered by Artificial Intelligence,” in Proc. IEEE Conf. Publication, 2024.
[3]. “Personalized Adaptive Learning Technologies Based on Machine Learning Techniques to Identify Learning Styles: A Systematic Literature Review,” IEEE Access, 2023.
[4]. “Toward Personalized Adaptive Learning Using Artificial Intelligence and Physiological Signals,” IEEE Journal, 2025.
[5]. “A Personalized Adaptive E-Learning System Based on Learner's Feedback and Learner's Multiple Intelligences,” in Proc. IEEE Conf. Publication, 2015.
[6]. “AI-Based Adaptive Personalized Content Presentation and Exercises Navigation for an Effective and Engaging E-Learning Platform (APPEAL),” Educ. Inf. Technol., 2022.
[7]. “Artificial Intelligence in Adaptive Education: A Systematic Review of Techniques for Personalized Learning,” Discover Educ., 2025.
[8]. “Personalized Artificial-Intelligence-Enhanced Learning: Issues, Challenges, and Solutions,” Int. Res. J. Modernization Eng., Technol. Sci., vol. 5, no. 12, 2023.
[9]. “Personalized Learning and Adaptive Systems: AI-Driven Educational Innovation and Student Outcome Enhancement,” Int. J. Educ. Humanities, 2025.
[10]. G. J. Hwang, H. Xie, B. W. Wah, and D. Gašević, “Vision, challenges, roles and research issues of Artificial Intelligence in Education,” Comput. Educ.: Artif. Intell., vol. 1, p. 100001, 2020.
[11]. A. Bhutoria, “Personalized education and artificial intelligence in the United States, China, and India: A systematic review using a human-in-the-loop model,” Comput. Educ.: Artif. Intell., vol. 3, p. 100068, 2022.
[12]. W. Holmes, M. Bialik, and C. Fadel, Artificial Intelligence in Education: Promises and Implications for Teaching and Learning. Boston, MA: Center for Curriculum Redesign, 2019.
[13]. C. Diwan, S. Srinivasa, G. Suri, S. Agarwal, and P. Ram, “AI-based learning content generation and learning pathway augmentation to increase learner engagement,” Comput. Educ.: Artif. Intell., vol. 4, pp. 100-110, 2023.
[14]. R. Berondo, “Harnessing the power of artificial intelligence for personalized learning in education,” Eur. Chem. Bull., vol. 12, no. 10, pp. 1243-1251, 2023.
[15]. R. S. J. d. Baker and P. S. Inventado, “Educational data mining and learning analytics,” in Learning Analytics, J. A. Larusson and B. White, Eds. New York, NY: Springer, 2014, pp. 61-75.

