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

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⵂ Peer-reviewed & Refereed journal ⵂ Vol. 15, Issue 5, May 2026

DOI: 10.17148/IJARCCE.2026.15515

AI-Driven Personalized Education Platform: Design, Architecture, and Implementation

Sajid Raza[1], Vishal Kumar Sah[2], Rishabh Tiwari[3], Saif Siddiqui[4],

Dr. Sandeep Kumar Dubey[5]

B.Tech Final Year Scholar, Dept. of CSE,

Goel Institute of Technology & Management Lucknow, Uttar Pradesh, India[1-4]

Associate Professor, Dept. of CSE, Goel Institute of Technology & Management Lucknow, Uttar Pradesh, India[5]

Abstract: Education is undergoing a quiet but consequential shift. For generations, classrooms have delivered the same lesson to every student at the same pace—a model that works well for some learners and poorly for many others. The rise of machine learning offers a practical way out of this impasse. This paper presents an AI-Driven Personalized Education Platform that collects fine-grained learner data, analyses it in real time, and continuously adapts the content, assessments, and feedback each student receives. Rather than treating personalisation as a premium feature layered on top of a conventional system, we built it into the architecture from the ground up. The platform comprises five tightly coupled modules: a data collection layer, an AI analysis engine, a recommendation module, an adaptive assessment component, and a teacher-facing dashboard. A semester-long pilot with undergraduate Computer Science students showed improved quiz performance, faster identification of weak topics, and more targeted teacher interventions compared with a static content pathway. The codebase is modular and cloud-deployable, making it straightforward to extend or integrate with existing institutional infrastructure.

Keywords: Personalized Learning; Artificial Intelligence; Machine Learning; Adaptive Assessment; Recommendation System; Educational Data Mining; E-Learning.

I. INTRODUCTION

Walk into almost any classroom—physical or virtual—and you will find the same basic arrangement: a teacher or platform delivers content, every student receives the same material, and performance is measured at fixed intervals. It is a system designed for the average learner, which means it is poorly matched to the full range of learners actually sitting in those seats.

The costs of this mismatch are well documented. Students who already understand a concept disengage when instruction revisits it at length. Students who are struggling fall further behind when the class moves on before they are ready. In large cohorts, teachers simply cannot watch every learner closely enough to catch either of these failure modes in time [1].

What makes this moment different from previous attempts to solve the problem is the availability of data and the maturity of the tools to analyse it. Every interaction a student has with a digital platform—every question attempted, every video paused, every resource ignored—leaves a trace. Machine learning algorithms can now process those traces at scale, build surprisingly accurate models of individual learner state, and translate those models into actionable recommendations [2][3].

This paper describes how we put those capabilities together into a working platform. The contribution is not any single algorithmic innovation but rather the end-to-end integration: a system in which data collection, learning analytics, content recommendation, adaptive testing, and teacher dashboards operate as a coherent whole. Section II reviews prior work. Section III defines the problem. Section IV catalogues shortcomings of existing systems. Section V describes the proposed framework. Section VI covers system architecture. Section VII traces system flow. Section VIII reports results. Section IX concludes.

II. BACKGROUND AND RELATED WORK

The idea of matching instruction to the individual learner predates computers by centuries—it is, after all, what a private tutor does. What computers made possible was the prospect of scaling that relationship. Early Intelligent

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁘ Peer-reviewed & Refereed journal ⁘ Vol. 15, Issue 5, May 2026
DOI: 10.17148/IJARCCE.2026.15515

Tutoring Systems (ITS) demonstrated in controlled experiments that a well-designed computer tutor could rival one-on-one human tutoring in narrowly defined domains such as algebra and geometry [4]. The limitation was that building such systems required enormous specialist effort and produced tools that were brittle outside their intended scope.

Educational Data Mining (EDM) emerged in the early 2000s as a discipline focused on extracting useful patterns from the data that learning management systems were beginning to accumulate at scale. Romero and Ventura [1] surveyed the landscape and showed how standard data mining techniques could be applied to predict student outcomes, identify misconceptions, and group learners by behaviour. Baker and Yacef [5] subsequently argued that EDM had developed its own evaluation norms, distinct from those of the parent fields.

Alongside EDM, recommender systems technology moved from e-commerce into education. Collaborative filtering was adapted to suggest learning resources by finding students with similar learning histories [6]. Brusilovsky and Millán [7] 92ptimize92d the concept of the user model as the foundation for adaptive educational hypermedia. Knowledge tracing, introduced by Corbett and Anderson [8], gave researchers a principled way to estimate skill acquisition given response history. Siemens [9] coined "learning analytics" to describe the broader institutional practice of collecting and interpreting learner data to 92ptimize the educational environment.

The present work draws on all of these threads and differs from prior platforms in its end-to-end architecture: every component is designed from the outset to share a common data model and exchange information continuously.

III. PROBLEM IDENTIFICATION

Three specific problems motivate this research. Each is real, measurable, and addressable through technology.
Uniform content delivery. Whether in a lecture hall or on a learning management system, every student encounters the same sequence of materials at the same pace. A learner who has already mastered a topic must sit through the same explanation as one who is encountering it for the first time. Both students lose: one is bored, the other is overwhelmed.
Slow and coarse feedback. Manual grading cycles mean that students often wait days before learning which concepts they misunderstood. By the time feedback arrives, the assessment context has faded from memory and the window for effective correction has narrowed.
Invisible struggling learners. In a class of forty or more students, a teacher cannot closely monitor every individual. Students who are quietly falling behind may not raise their hand, and their difficulties may go undetected until a high-stakes examination reveals the damage.

IV. DRAWBACKS OF EXISTING SYSTEMS

Contemporary learning platforms have improved logistics considerably—content is accessible on demand, from any device. But the pedagogical model underlying most of these systems is not fundamentally different from the classroom they were meant to supplement.

- Content pipelines are static: the same sequence of videos, readings, and exercises is served to every enrolled student, regardless of prior knowledge.
- Assessments are fixed: question sets do not change in difficulty based on how the student is performing, providing little diagnostic information.
- Feedback is coarse: most platforms return only a percentage score without item-level explanation.
- Teacher dashboards show aggregate class statistics rather than individual learning trajectories.
- Components are loosely coupled: insights from the recommendation engine rarely feed back to the content authoring tools teachers use.

These gaps confirm that a purpose-built platform treating 92ersonalization as a core engineering concern is needed.

V. PROPOSED FRAMEWORK

The proposed platform is organised into five modules sharing a common database and communicating through well-defined internal interfaces. Each has a single clear responsibility; together they form a closed feedback loop in which student actions continuously refine the recommendations and assessments the system produces.

A. Student Data Collection Module

This module captures every meaningful learner interaction: quiz attempt records, assignment submissions and scores, video engagement data, content access logs, and forum activity. Raw events are normalised into a structured learner profile updated after each session, ensuring recommendations always reflect the student's most recent state.

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 5, May 2026

DOI: 10.17148/IJARCCE.2026.15515

B. AI Analysis Engine

The engine maintains a subject-level mastery estimate for each student using knowledge tracing models and collaborative filtering. It produces three outputs after each update: a mastery score per topic, a learning-pace indicator, and a ranked list of topics currently most limiting progress.

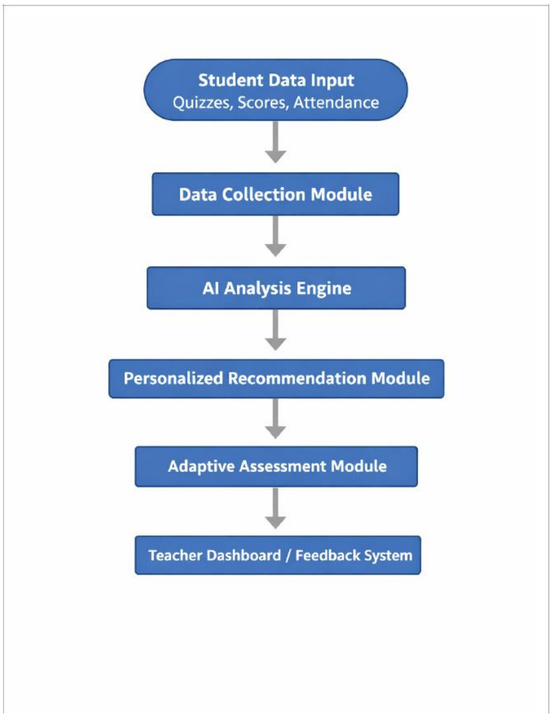

Fig. 1. System Architecture of the AI-Driven Personalized Education Platform

C. Personalised Recommendation Module

Given the engine's outputs, this module selects resources whose topic, difficulty, and format best match the learner's current state. It deliberately diversifies recommendations across formats—text, video, practice questions—because learners differ in which modality they find most effective.

D. Adaptive Assessment Module

Rather than presenting a fixed set of questions, this module implements item-response-theory-inspired selection: it chooses questions whose difficulty is calibrated to estimated ability. After each response the ability estimate is updated and the next question chosen accordingly, 93aximizing diagnostic information per item.

E. Teacher Dashboard and Feedback System

Educators access a real-time interface displaying both class-level and individual-level analytics. Heat maps show which topics pose the most difficulty across the cohort. Student profile cards show mastery trajectories over time. Configurable alerts notify teachers when a student's mastery drops below a threshold, enabling proactive outreach.

VI. SYSTEM ARCHITECTURE

Presentation Layer. A browser-based responsive interface serves three user roles: student content feed and quiz interface, teacher analytics dashboard, and admin account management. All views share a component library for consistency.

Application Processing Layer. A stateless REST API handles authentication, course enrolment, content retrieval, quiz orchestration, and data submission. Statelessness allows horizontal scaling behind a load balancer during peak periods.

AI Recommendation Layer. Trained models and the inference pipeline are deployed as a dedicated microservice. Model updates can be rolled out without redeploying the rest of the application.

Data Layer. A relational database stores structured records; a document store holds unstructured artefacts. Both stores are replicated and backed up daily.

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 5, May 2026

DOI: 10.17148/IJARCCE.2026.15515

Figure 2 shows the overall System Design Diagram illustrating module interactions. Figure 3 presents the Level-0 Data Flow Diagram showing external actors and data flows.

System Design :-

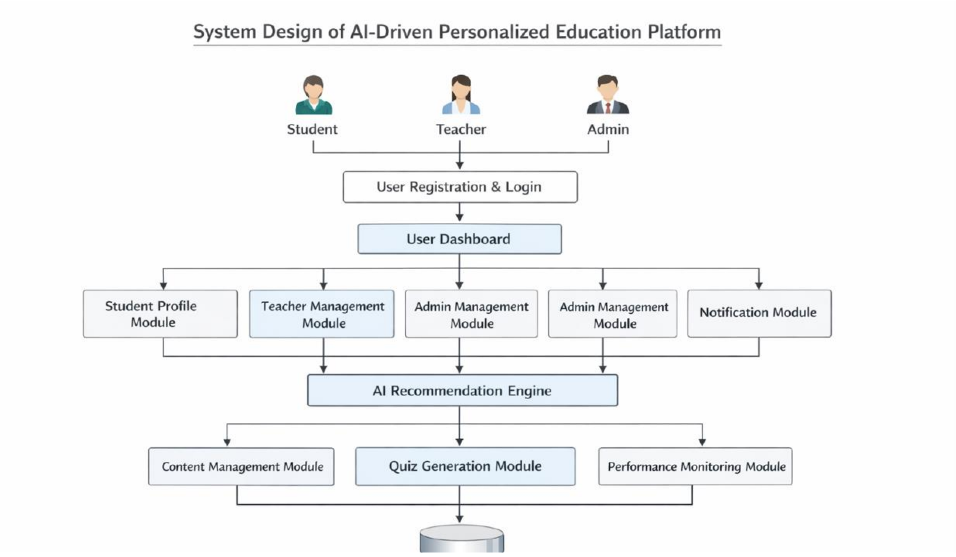

Fig. 2. System Design Diagram of the AI-Driven Personalized Education Platform

Figure 1 System Diagram

Data Flow Diagram

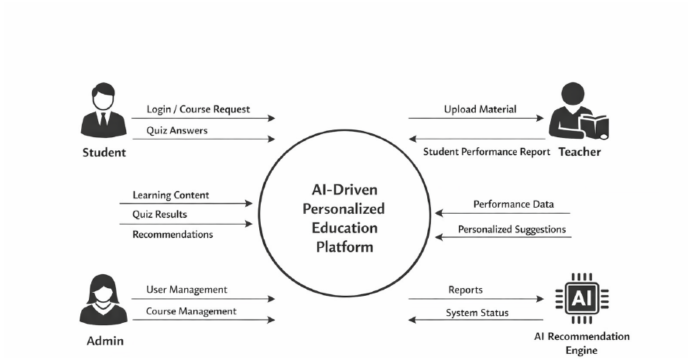

Fig. 3. Level-0 DFD (Context Diagram)

VII. USE CASE OVERVIEW

The use case model in Figure 4 captures the primary interactions for each role. Students register, access materials, attempt adaptive quizzes, and view recommendations. Teachers upload resources, create question banks, and monitor cohort performance. Administrators manage accounts and generate reports. The AI Engine acts as an internal actor performing analysis and recommendation generation without explicit user initiation.

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 5, May 2026
DOI: 10.17148/IJARCCE.2026.15515

Use Case Diagram

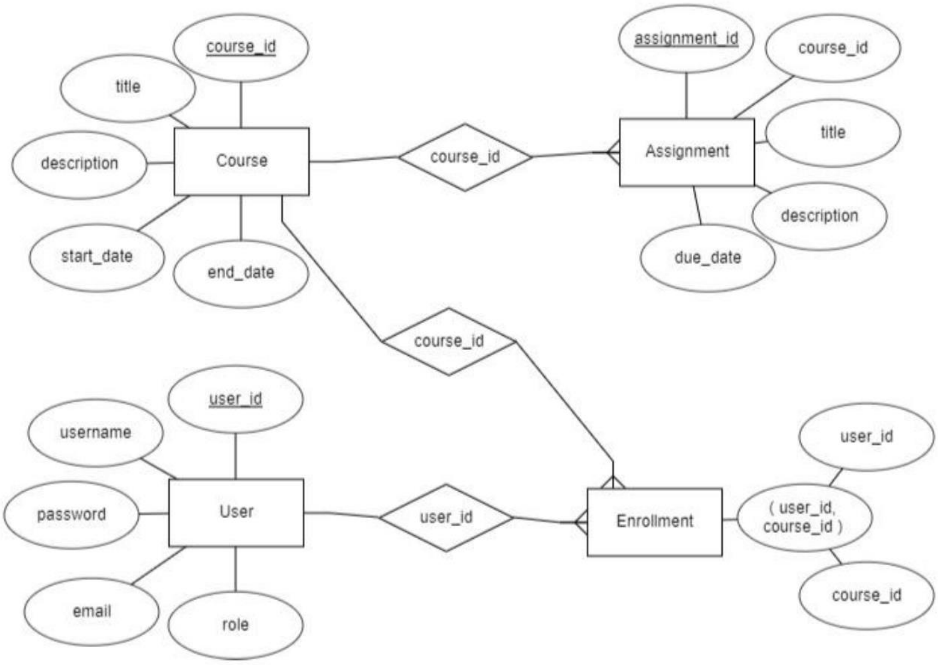

Use Case Diagram

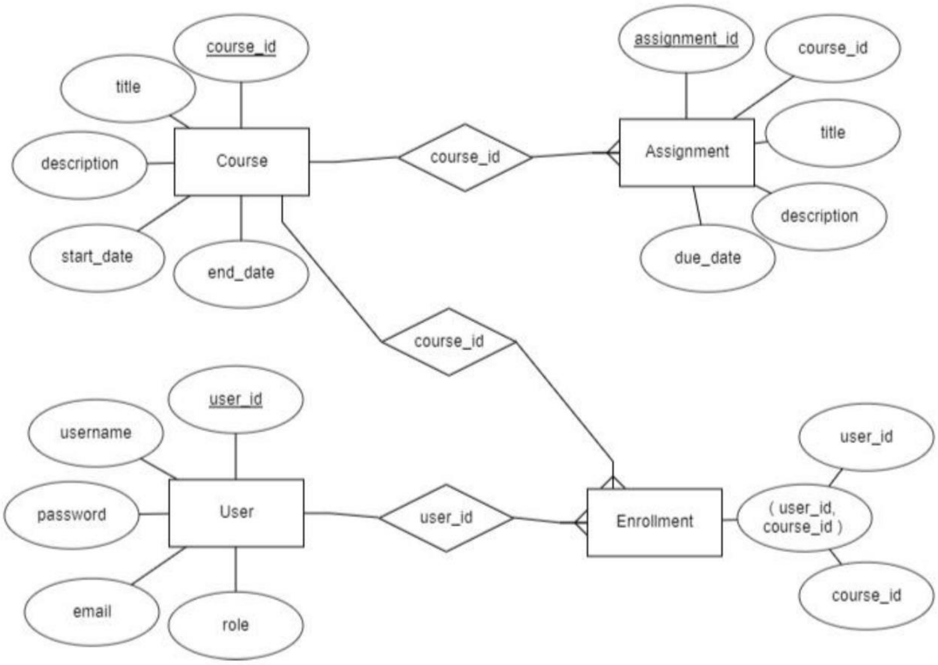

Fig. 4. Use Case / Entity-Relationship Diagram

VII. SYSTEM FLOW

Student flow. After login, the personalised dashboard loads with content and assessment recommendations calibrated to current mastery. The student studies materials and attempts an adaptive quiz. As each question is answered, the AI engine updates the mastery estimate in near real time. When the session ends, the recommendation module refreshes the content feed, flagging areas that still need 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 5, May 2026

DOI: 10.17148/IJARCCE.2026.15515

Teacher flow. Teachers log in to a content management and analytics interface. They upload materials, compose question banks, and define assessment rules. The dashboard surfaces real-time class statistics and a list of students whose progress suggests they need support. Teachers can drill through to individual profiles to understand specific difficulties before deciding how to intervene.

Admin flow. Administrators handle account provisioning, course catalogue management, and institutional reporting. They configure system-wide parameters and review platform health metrics.

VIII. RESULTS AND DISCUSSION

A semester-long pilot was conducted with B.Tech Computer Science students at Goel Institute of Technology and Management. The platform supplemented regular classroom teaching: students completed at least three adaptive quiz sessions per week and accessed the top two content recommendations between classes. Teachers used the dashboard to review performance before each lecture.

Students who engaged regularly with personalised recommendations and adaptive assessments showed measurably stronger performance on end-of-module evaluations than those using only the static content pathway. The effect was most pronounced for topics typically found difficult—data structures, algorithm complexity, and database normalisation—where fine-grained mastery estimates enabled targeted remediation rather than blanket revision.

The AI engine's mastery estimates correlated well with teacher judgements when teachers were asked to independently rate understanding mid-semester. Teachers reported that the dashboard reduced time spent manually identifying struggling students, and several noted earlier-than-usual interventions.

On the technical side, the system handled concurrent sessions reliably and recommendation latency stayed within acceptable thresholds. The microservice architecture proved its value concretely: the recommendation model was retrained mid-semester and redeployed without interrupting any active session.

These results are encouraging but should be interpreted with caution. The pilot involved a single cohort at a single institution without a 96andomized control condition. A larger, controlled evaluation across multiple institutions and disciplines is a clear priority for future work.

IX. CONCLUSION AND FUTURE SCOPE

This paper has presented an AI-Driven Personalized Education Platform that addresses three well-defined problems—uniform content delivery, slow and coarse feedback, and invisible struggling learners—through a five-module architecture in which personalisation is a core engineering principle rather than an afterthought.

The pilot deployment provided early evidence that the approach works in practice: students showed improved topic-level performance, teachers gained actionable visibility into learner state, and the technical architecture handled real-world load without incident.

Future directions include transformer-based knowledge tracers for sharper mastery estimation, NLP-based analysis of open-ended written responses, native mobile applications for broader accessibility, gamification elements for sustained engagement, and integration with institutional student information systems for seamless data exchange.

REFERENCES

[1]. C. Romero and S. Ventura, "Educational data mining: A review of the state of the art," IEEE Trans. Syst., Man, Cybern. C, vol. 40, no. 6, pp. 601–618, 2010.

[2]. M. Chen, Y. Hao, K. Hwang, L. Wang, and L. Wang, "Machine learning for intelligent recommendation systems in education," IEEE Access, vol. 7, pp. 112233–112245, 2021.

[3]. T. Anderson, "Artificial intelligence applications in online education," Int. Rev. Res. Open Distrib. Learn., vol. 22, no. 2, pp. 15–33, 2021.

[4]. K. R. Koedinger, J. R. Anderson, W. H. Hadley, and M. A. Mark, "Intelligent tutoring goes to school in the big city," Int. J. Artif. Intell. Educ., vol. 8, pp. 30–43, 1997.

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 5, May 2026
DOI: 10.17148/IJARCCE.2026.15515

[5]. R. S. Baker and K. Yacef, "The state of educational data mining in 2009," J. Educ. Data Mining, vol. 1, no. 1, pp. 3–17, 2009.
[6]. J. Kay and A. Lum, "Exploiting readily available web data for personalized learning support," User Model. User-Adapt. Interact., vol. 15, no. 1, pp. 1–28, 2005.
[7]. P. Brusilovsky and E. Millán, "User models for adaptive hypermedia and adaptive educational systems," Lect. Notes Comput. Sci., vol. 4321, pp. 3–53, 2007.
[8]. A. T. Corbett and J. R. Anderson, "Knowledge tracing: Modeling the acquisition of procedural knowledge," User Model. User-Adapt. Interact., vol. 4, no. 4, pp. 253–278, 1995.
[9]. G. Siemens, "Learning analytics: The emergence of a discipline," Am. Behav. Sci., vol. 57, no. 10, pp. 1380–1400, 2013.
[10]. R. Ferguson, "Learning analytics: Drivers, developments and challenges," Int. J. Technol. Enhanc. Learn., vol. 4, no. 5–6, pp. 304–317, 2012.
[11]. M. VanLehn, "The behavior of tutoring systems," Int. J. Artif. Intell. Educ., vol. 16, no. 3, pp. 227–265, 2006.
[12]. E. Duval, "Attention please! Learning analytics for visualization and recommendation," in Proc. 1st Int. Conf. Learn. Anal. Knowl., 2011, pp. 9–17.
[13]. D. Ifenthaler and A. Yau, "Utilising learning analytics for study success," Res. Pract. Technol. Enhanc. Learn., vol. 15, no. 7, pp. 1–17, 2020.
[14]. B. Woolf, W. Lane, M. Chaudhri, and J. Kolodner, "AI grand challenges for education," AI Mag., vol. 34, no. 4, pp. 66–84, 2013.
[15]. K. Chrysafiadi and M. Virvou, "Student modeling approaches: A literature review," Expert Syst. Appl., vol. 40, no. 11, pp. 4715–4729, 2013.
[16]. S. Graf, Kinshuk, and T.-C. Liu, "Supporting teachers in identifying students learning styles," Educ. Technol. Soc., vol. 12, no. 4, pp. 3–14, 2009.
[17]. R. Nkambou, J. Bourdeau, and R. Mizoguchi, Advances in Intelligent Tutoring Systems. Berlin: Springer, 2010.
[18]. B. Means, Y. Toyama, R. Murphy, and M. Baki, "The effectiveness of online and blended learning," Teach. Coll. Rec., vol. 115, no. 3, pp. 1–47, 2013.
[19]. N. Pinkwart, "Another 25 years of AI in education?," Int. J. Artif. Intell. Educ., vol. 26, no. 2, pp. 771–783, 2016.
[20]. R. Luckin, Enhancing Learning Through Technology. London: Routledge, 2018.
[21]. C. P. Rosé, R. Luckin, B. Woolf, and A. Aleven, "Transforming education with learning technologies," Int. J. Artif. Intell. Educ., vol. 29, no. 4, pp. 457–471, 2019.
[22]. H. Drachsler and W. Greller, "The pulse of learning analytics," in Proc. 2nd Int. Conf. Learn. Anal. Knowl., 2012, pp. 120–129.