International Journal of Sciences and Innovation Engineering
(Peer-Reviewed, Open Access, Fully Refereed International Journal)
website: www.ijsci.com
Volume 3 Issue 8 , August 2026
E-ISSN : 3049-0251
DOI: https://doi.org/10.70849/ijsci03082611087

A GenAI-Based Adaptive Tutoring and Intelligent Assessment Framework for Personalized Learning

Prof. Aruna M G[1], Dhyan Gowda[2], Priya Ramesh Prasad[3], Punith Kumar[4], Sachin N[5]

[1,2,3,4,5] *Dept. of AI & ML Dayananda Sagar University Bengaluru, India.*

Article Info

Article History:

Published: 05 August 2026

Publication Issue:
Volume 3, Issue 8
August-2026

Page Number:
43-48

Corresponding Author:
Prof. Aruna M G

Abstract:

Modern academic institutions face a fundamental gap: instruction is designed for the average student, leaving individuals with specific weaknesses without any targeted support mechanism. Standardised course pipelines and identical assessments for all students have consistently failed to close this instructional gap. Advances across Artificial Intelligence (AI), Natural Language Processing (NLP), and Generative AI now make it feasible to construct learning environments that actively evolves as each student progresses [1].This paper introduces EduMind, a unified tutoring and assessment platform designed around a dual-track evaluation model. Closed-form questions are evaluated using fixed-logic scoring for consistent results, while open-ended answers are assessed by computing meaning-level correspondence with expert reference responses.The combined output enables fine-grained identification of both proficient and deficient knowledge areas at the topic level [2][3].EduMind's embedded Generative AI module transforms evaluation data directly into targeted instructional content, addressing identified gaps with structured explanations and study material packaged into downloadable PDF reports. The system demonstrates how assessment and tutoring can be unified into a seamless workflow, and remained operationally stable throughout all testing phases.

Keywords: Generative AI, adaptive learning, intelligent tutoring systems, automated assessment, personalized feedback, NLP

1. INTRODUCTION

Science and engineering curricula have witnessed substantial digitization over the past decade, yet a core pedagogical problem persists: students rarely receive feedback that is both immediate and calibrated to their specific knowledge gaps. Traditional course delivery is fundamentally uniform: every student receives the same content at the same speed, judged by a single shared rubric that cannot account for where each individual actually stands in their understanding. As a result, misunderstandings quietly build up over time — neither the student nor the teacher realises how serious these gaps have become until a formal exam brings them to light.

Widely deployed course management platforms address logistical coordination effectively but offer little in the way of dynamic pedagogical response. Their evaluation components are largely static, scoring answers against fixed keys and returning marks without interpretation.

Students who score poorly receive no structural guidance about which concepts failed them or how to address the gap.The emergence of AI and NLP as practical engineering tools has changed the design space for educational software considerably. Using these tools, a system can analyse the actual meaning behind a student's answer, generate explanations suited to a specific knowledge gap, and update its model of each learner as more response data becomes available [1][7].

Generative AI extends this further by enabling the real-time synthesis of explanations tuned to a learner's demonstrated level of understanding rather than pulling from a pre-built library [5] [7].

Despite this progress, In most deployed educational platforms, evaluation, feedback, and tutoring operate as completely independent subsystems — a student's performance data gathered in one stage simply does not flow into the next. A student's performance on an assessment does not automatically reconfigure the tutoring support they receive next. EduMind is designed specifically to close this gap. It couples an intelligent evaluation engine with a Generative AI tutoring module so that assessment outcomes directly govern what explanatory content is produced for the student [5][7].

The system additionally synthesizes this content into structured, student-specific study reports available for independent revision. This paper describes the architecture, implementation choices, and testing outcomes that characterize EduMind as a practical step toward truly adaptive educational technology [2][3].

2. SURVEY METHODOLOGY

The literature review was structured around three primary research questions: how AI-based techniques can be applied to individualize student learning pathways; which evaluation methods are best suited to both closed-form and open-ended student responses; and what role Generative AI can play in producing pedagogically sound explanatory content on demand [5][7].

A secondary thread explored the relationship between granular topic-level performance diagnostics and measurable learning improvement. Source material was gathered from IEEE Xplore, Springer, Elsevier, and arXiv preprints.

Inclusion criteria prioritized works with concrete system implementations involving AI, machine learning, or NLP in an educational context---particularly intelligent tutoring systems, automated response evaluation, and adaptive feedback mechanisms. .Purely theoretical contributions or those outside the educational domain were excluded from primary analysis.

The selected papers were examined for recurring methodological patterns: approaches to semantic response evaluation, rule-based versus learning-based scoring, adaptive content delivery, and the integration of Generative AI for explanation synthesis [2][3].

The most consistent finding across the surveyed literature was the architectural separation of evaluation and tutoring as independent subsystems with limited inter-component communication---a limitation that directly motivated the integrated design of EduMind.

Table 1 : Paper Reference

Database	Total	Primary relevant	Most relevant
IEEE Xplore	12	12	10
Digital Library	5	4	3
Springer	5	4	4
Elsevier	4	3	3
MDPI arXiv	4	3	4
Total	30	26	24

3. SYSTEM ARCHITECTURE

EduMind is structured as a modular pipeline in which assessment data flows continuously into tutoring logic, producing personalized instructional content without requiring manual intervention between stages. Each module has a well-defined interface, enabling independent development and future integration with third-party educational platforms [1][6].

The student-facing entry point is a web interface where users authenticate, choose a subject domain, and engage with quiz sessions. The assessment module controls the sequencing and presentation of questions, collects responses, and maintains session state throughout the interaction.

Internally, two parallel evaluation paths operate depending on question format. A dedicated question controller manages progression logic.

The scoring subsystem applies deterministic rules to objective responses, producing consistent and auditable results. For free-text descriptive answers, a semantic matching engine computes similarity against curated reference responses, capturing conceptual accuracy even when student phrasing diverges from the expected formulation [2][3].

Topic-level aggregation runs following individual question scoring. The system accumulates performance signals across all questions within a topic and flags those falling below a defined threshold for remediation. This targeted flagging prevents the tutoring module from generating content indiscriminately across the full subject scope.

Fig 1. Workflow of the proposed EduMind system for adaptive assessment and evaluation. Illustrates the integration of automated scoring, AI tutoring, and personalized report generation.

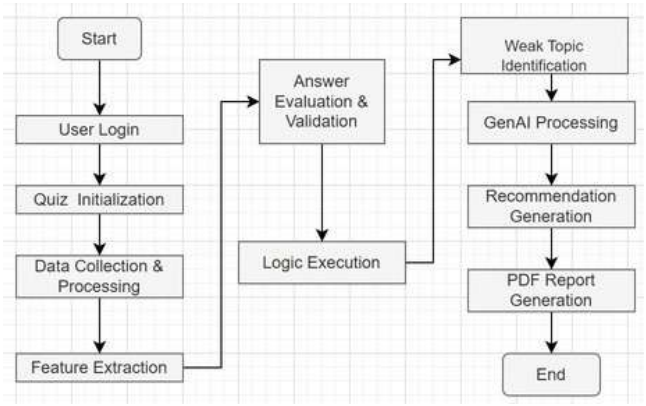

As shown in Fig.1, the system incorporates multiple supporting components that operate in coordination during evaluation. The question controller oversees the organization and progression of questions, while the scoring mechanism assesses objective responses using pre defined logic. For descriptive answers, similarity based techniques are applied to match student responses against reference answers, enabling the detection of conceptual accuracy independent of surface phrasing [2][3].

The tutoring module receives the list of flagged topics and constructs explanatory content through its Generative AI interface. The resulting material is calibrated to the specific conceptual gaps indicated by the assessment data rather than recycling generic subject matter [5][7].

A report generation module assembles this content into a structured PDF document containing topic-specific explanations and revision guidance. The document is made available to the student immediately upon session completion, supporting self-directed study outside the platform. The modular design ensures that any individual component can be replaced or upgraded without disrupting the overall data flow.

4. SOFTWARE IMPLEMENTATION

EduMind was realized as a browser-accessible web application, integrating adaptive evaluation and AI-driven tutoring within a unified platform. Implementation decisions prioritized low-latency data transfer between modules so that learning feedback could reach students promptly at the close of each assessment s-ession.

The student interface provides authentication, subject selection, and quiz navigation.Interface complexity was deliberately minimized to reduce cognitive overhead and focus the user experience on the assessment content itself.

All response data collected through the interface is transmitted to the backend processing layer for evaluation.

The backend is implemented on a Python-based server framework that orchestrates the execution of the core business logic. It coordinates quiz delivery, persists response records, and manages inter-module communication throughout the session lifecycle. This centralized architecture ensures consistent data handling regardless of concurrent session load.

EduMind evaluates two different question types through separate processing paths. For closed-form items such as multiple choice, a logic engine checks the response against stored correct answers and produces a consistent, reproducible score regardless of minor phrasing differences.Free-text answers go through a meaning-extraction pipeline: the student's response is encoded and scored against an expert-written reference answer, capturing conceptual correctness even when the vocabulary used differs considerably from the expected phrasing [2][3].

After all questions are evaluated, EduMind groups scores by topic to produce a subject-level summary that identifies where the student's knowledge is solid and where understanding is incomplete. Only topics that fall below the remediation threshold are forwarded to the tutoring module.

For each flagged topic, EduMind builds a prompt containing the topic context and the student's specific performance signal, then submits it to the generative model backend. The response is formatted as step-by-step explanations with worked examples — directly addressing the point of confusion rather than providing a generic subject overview [5][7].

The report generation component receives this content and formats it into a downloadable PDF, organizing explanations by topic alongside actionable study suggestions. Students retain this document for use during independent revision after the session ends.

End-to-end testing across all modules confirmed stable execution. The system processed interactions, completed evaluations, and produced reports without encountering workflow failures, indicating reliability suitable for deployment in live academic settings.

5. INTELLIGENT ASSESSMENT AND AUTOMATED EVALUATION TECHNIQUES

Automated evaluation has moved well beyond simple answer-matching as AI and NLP capabilities have matured. Early educational software measured correctness through character-level string comparisons, a method that penalizes valid responses expressed in alternative but semantically equivalent ways, and that cannot detect partial understanding at all.

Contemporary approaches instead operate at the level of meaning, measuring whether a student's response reflects genuine comprehension of the underlying concept [2][3].

The result is an evaluation layer capable of distinguishing students who understand a concept but express it differently

from those who do not understand it at all. This granularity has direct value for remediation targeting.

Beyond per-question scoring, the system performs a second-order analysis at the topic level. By aggregating scores across all questions within a topic, EduMind constructs a performance map that distinguishes domains of strength from domains requiring intervention. Tutoring resources are allocated only to the latter, ensuring that the support offered is directly proportionate to the student's actual needs rather than uniformly distributed [1][6].

A critical design principle underlying this approach is the tight coupling between evaluation and remediation. In architectures where these functions are separated, the time lag between identifying a knowledge gap and delivering content to address it degrades the pedagogical value of the feedback. EduMind eliminates this lag by making the evaluation output the direct input to the tutoring module, transforming assessment from a terminal event into the initiating step of a learning cycle.

6. GENERATIVE AI IN PERSONALIZED LEARNING SYSTEMS

The practical availability of large-scale generative models has introduced a qualitatively different capability into educational technology: the ability to compose instructional material on demand, tailored to what a specific student revealed during their session, rather than drawing from a fixed catalogue of pre-written content. This shift has implications for scalability, because the system is no longer limited by the breadth of content that was prepared in

advance, and for relevance, because the generated material can address the exact conceptual gaps a specific student has just demonstrated [5][7].

Experimental work in this area has established that generative models can decompose complex subject matter into accessible explanations when prompted appropriately, and that learners benefit more from correction focused precisely on the error they made than from a broad review of the general topic area [5][7].

The prerequisite for this outcome is a reliable mechanism for identifying what the individual student actually needs---which is precisely what EduMind's evaluation layer provides.

Within EduMind, Generative AI serves as the execution engine of the tutoring module. Once weak topics are flagged from assessment data, the module constructs prompts that provide the model with the relevant topic context and the observed performance signal, instructing it to produce an explanation that targets the specific conceptual failure.The output is integrated directly into the student's report rather than presented as a separate interaction, preserving session continuity.

Operational challenges associated with Generative AI in this context include output consistency and depth. Without carefully engineered prompts and output validation, generated explanations can vary significantly in quality across sessions. EduMind addresses this through structured prompt templates that constrain the model's output format and anchor it to the specific topic under remediation, reducing variance in explanation quality.

The key architectural distinction in EduMind is that Generative

AI is not a supplementary feature appended to an otherwise conventional system. It is a core pipeline stage whose inputs are determined by assessment outcomes and whose outputs constitute the primary tutoring artifact the student receives. This integration defines the system's adaptive character.

7. SYSTEM PERFORMANCE AND ANALYSIS

EduMind was evaluated across its complete operational workflow to assess correctness, stability, and the quality of adaptive support produced. Testing confirmed that both evaluation tracks---rule-based objective scoring and semantic similarity-based descriptive scoring---operated accurately and produced consistent results across repeated sessions with the same inputs.

Topic-level performance profiling successfully partitioned assessed topics into remediation candidates and satisfactorily-understood domains based on accumulated scoring data. The semantic similarity engine correctly resolved cases where students expressed equivalent concepts using different vocabulary, assigning credit appropriately rather than penalizing surface-level variation [2][3].

Tutoring content generated for flagged topics was coherent and topically focused. Students received explanations addressing only the specific concepts that their assessment performance had identified as weak, rather than general topic summaries.

The report assembly module produced structurally consistent PDF documents in all test cases, with explanations and study guidance organized by topic in a format suitable for independent revision use.

No workflow failures or significant latency issues were observed during testing. The system executed the full sequence---session management, response evaluation, topic profiling, content generation, and report production---without interruption, suggesting the architecture is sufficiently robust for deployment in real academic environments [2][3].

8. CONCLUSION

This paper presented EduMind, a GenAI-integrated tutoring and assessment platform engineered to close the gap between performance measurement and personalized learning support. The system pairs rule-based and semantic evaluation methods to produce a nuanced picture of each student's knowledge state, then converts that picture into targeted instructional content through a tightly coupled Generative AI module.

The core contribution of this work is architectural: by eliminating the separation between the evaluation subsystem and the tutoring subsystem, EduMind ensures that the support a student receives is directly and specifically derived from what the assessment revealed about their understanding. This contrasts with the prevailing pattern in educational platforms, where feedback remains generic and disconnected from individual performance data.

Testing demonstrated operational stability across all pipeline stages and confirmed the system's ability to produce educationally useful content in response to real student performance signals. These outcomes position EduMind as a concrete and extensible contribution toward adaptive educatio-educational technology that genuinely responds to the individual learner.

References

1. H. Chen, X. Liu, and Y. Zhang, "AI-Powered Intelligent Tutoring Systems: A Review, " IEEE Access, vol. 9, pp. 125678–125689, 2021.
2. L. Sun, J. Li, and Y. Chen, "AI-Based Automated Feedback Systems for Research Manuscripts, " IEEE Transactions on Learning Technologies, vol. 15, no. 3, pp. 345–356, 2022.
3. P. Singh, G. Gupta, and S. Kumar, "Intelligent Feedback Engines for Student Reports Using Transformer Models, " Proceedings of the International Conference on Artificial Intelligence in Education, 2023, pp. 210–218.
4. B. Lee, S. Kim, and H. Park, "AI-Driven Research Paper Recommendation Systems, " Springer Journal of Information Retrieval, vol. 22, no. 4, pp. 455–468, 2019
5. F. Ahmed, M. Rahman, and K. Hasan, "Deep Learning for Automated Essay and Paper Grading, " IEEE Access, vol. 9, pp. 87543–87555, 2021.
6. A. Létourneau, P. Tremblay, and J. Roy, "A Comprehensive Review of AI-Based Intelligent Tutoring Systems, " arXiv preprint arXiv:2501.03421, 2025.
7. W. Villegas-Ch, X. García-Ortiz, and S. Sánchez, "Adaptive Intelligent Tutoring Systems for STEM Education, " Springer Journal of Educational Technology Research, vol. 18, no. 2, pp. 145–158, 2025.
8. A. W. Fodouop Kouam, "Effectiveness of Intelligent Tutoring Systems Across Programming Experience Levels, " SpringerLink Journal of Computer Science Education, vol. 32, pp. 85–96, 2024.