iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** IV **Month of publication:** April 2026

DOI: https://doi.org/10.22214/ijraset.2026.80701

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IV Apr 2026- Available at www.ijraset.com

LearnCurve: Enhancing Learning Outcomes Through AI-Driven Generation and Mastery-Based Progression

Dr. Vinay Kumar Nassa[1], Snehit Pandey[2]
HMR Institute of Technology and Management GGSIPU, Delhi 110036

ABSTRACT: ***Self-directed digital learning fails at the intersection of pacing and personalization: platforms deliver uniform content regardless of what individual learners already know or where their understanding breaks down. This paper presents LearnCurve, an AI-driven adaptive learning platform that addresses this through three integrated mechanisms: First an AI-based pipeline that generates structured, hierarchical learning roadmaps tailored to a learner's goal, skill level, and timeline. Second a Mastery Index (MI) algorithm that synthesizes six per-topic performance dimensions into a continuous score driving real-time content adaptation.Third structured focus sessions with AI-generated post-session summaries that reinforce knowledge consolidation immediately after study. Unlike platforms that apply AI narrowly to recommendation or chatbot assistance, LearnCurve embeds it within a closed feedback loop spanning generation, assessment, adaptation, and reinforcement ensuring every learner interaction reshapeswhat that learner encounters next.***

Keywords: Adaptive learning systems, Mastery Index, AI roadmap generation, fine-tuned language models, personalized education

I. INTRODUCTION

A. The Pacing Problem in Digital Learning

The democratization of online education has created a paradox: access to learning resources is nearly universal, yet meaningful learning outcomes remain limited for many self-directed learners. Completion rates for self-paced digital courses remain below 15% [1], and longitudinal studies consistently identify the absence of personalized pacing—rather than content quality—as a primary driver of disengagement [2]. Learners who already understand a concept are forced to revisit familiar material, while those lacking prerequisite knowledge struggle to absorb subsequent topics. Static curricula fail to accommodate either case.

The root of this issue is structural. Traditional course design assumes a uniform learner profile and delivers content in a fixed sequence. In reality, learners possess diverse prior knowledge, progress at different rates, and exhibit uneven mastery across topics. Bloom's Mastery Learning model (1968) demonstrated that most learners can achieve high proficiency given sufficient time and appropriately adapted instruction [3]. However, implementing such individualized instruction at scale requires continuous assessment and dynamic content adjustmenttasks that are impractical through manual methods alone.

B. AI as the Enabler of Scalable Personalization

Artificial intelligence provides a viable approach to overcoming this limitation. Modern language models can generate structured learning content on demand, reducing reliance on manually authored curricula. When combined with behavioural datasuch as quiz performance, response time, session activity, and self-reported confidencethese systems can support adaptive learning environments that adjust in real time to a learner's evolving knowledge state [4].

Despite these advances, most AI-enhanced learning platforms apply AI in a limited capacity, such as content recommendation or chatbot assistance. Few systems integrate AI across the entire learning lifecycle, including roadmap generation, assessment, adaptive scheduling, and post-session reinforcement. LearnCurve is designed to address this gap.

C. Research Contributions

This research makes the following contributions:

1. Adaptive Roadmap Engine:A structured AI-based generation pipeline that produces hierarchical learning roadmaps (Phases → Milestones → Topics) with estimated durations and resource suggestions, tailored to a specified learning goal and timeline.

©IJRASET: All Rights are Reserved

2. Mastery Index Algorithm:A multi-dimensional composite metric that integrates accuracy, time efficiency, completion ratio, attempt count, revision count, and confidence into a continuous mastery score (0–100), enabling dynamic classification of topics as Weak, Developing, or Mastered.
3. Focus Session Pipeline:A structured study workflow based on timed focus sessions, complemented by AI-assisted post-session summaries to support immediate knowledge consolidation.
4. Closed Adaptive Feedback Loop:An integrated system in which learner interactions continuously update mastery scores, which in turn dynamically adjust content sequencing and task prioritization without manual intervention.

II. LITERATURE REVIEW

A. Mastery Learning and Prerequisite-Based Progression

Bloom's Mastery Learning framework posits that the primary limitation of conventional instruction lies not in learner ability, but in uniform time allocation, where all learners are exposed to the same content for the same duration regardless of individual readiness [3]. The model emphasizes that learners should demonstrate competence in a topic before progressing to dependent material. Empirical studies across multiple domains report effect sizes ranging from 0.5 to 1.0 compared to traditional instruction [5].

Despite its effectiveness, implementing mastery-based progression at scale has remained challenging due to the need for individualized assessment and content adaptation. The emergence of AI-driven roadmap generation provides a scalable mechanism to operationalize mastery learning principles in digital environments.

B. Intelligent Tutoring and Adaptive Learning Systems

Intelligent Tutoring Systems (ITS), developed in the late 20th century, demonstrated that computational models of learner knowledge could achieve learning gains comparable to human tutoring [6]. Modern adaptive platforms, such as Carnegie Learning and ALEKS, extend these systems using techniques like Bayesian knowledge tracing and item response theory to estimate learner proficiency.

Empirical studies indicate that adaptive learning systems improve learning efficiency by 20–40% compared to static instruction [7]. However, these systems rely heavily on pre-authored content, structured item banks, and expert-designed knowledge graphs, limiting their scalability across diverse domains. AI-generated content offers a potential solution to reduce this dependency.

C. AI-Generated Educational Content

Recent advancements in transformer-based language models have significantly reduced the cost of generating structured educational content. Prior work demonstrates that schema-constrained AI systems can generate personalized learning pathways comparable in quality to those designed by domain experts, particularly in terms of coherence and sequencing [8]. Similarly, studies show that AI-generated assessment items can match human-authored questions in clarity and difficulty calibration [9].

A key challenge, however, lies in ensuring structural reliability. Educational systems require consistent, machine-parse able outputs rather than free-form text. This necessitates the use of constrained generation techniques and validation pipelines to ensure predictable and usable outputs.

D. Cognitive Load and Spaced Practice

Cognitive Load Theory (Sweller, 1988) establishes that working memory is limited, and that instructional design should minimize extraneous cognitive load while maximizing germane load to support effective learning [10]. Structured learning approaches, such as time-boxed study sessions, reduce decision-making overhead and improve focus.

Additionally, research on spaced practice demonstrates that distributed learning sessions with optimally timed review intervals can improve retention by 10–50% compared to massed learning [11]. These findings support the use of structured focus sessions and adaptive review scheduling in learning systems.

E. Formative Assessment and Feedback Loops

Formative assessment, defined as assessment used to guide subsequent learning, has been shown to significantly improve learning outcomes, with reported effect sizes between 0.4 and 0.7 [12]. A critical requirement is that assessment results must actively influence the learning process.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

Table 1: - Comparative Analysis of Adaptive Learning and AI-Generated Roadmap Research

S.No.	Author(s) and Year	Technology and Methodology	Results	Advantages/ Applications	Disadvantages/ Pitfalls
1	Bloom, 1968 [3]	Mastery Learning framework	Near-universal mastery achievable with pacing individualization	Theoretical basis for MI gating	Manual pacing impractical at scale
2	VanLehn, 2011 [6]	ITS effectiveness meta-analysis	ITS produces effect sizes approaching human tutoring (d=0.79)	Validates adaptive knowledge modelling	Requires expensive pre-authored item banks
3	Pane et al., 2017 [7]	Adaptive learning platform RCT	21% learning efficiency gain over static instruction	Validates adaptive scheduling approach	Lacks AI content generation
4	Martin & Sumithra, 2025 [8]	AI roadmap generation quality	Schema-constrained AI generates expert-comparable pathways	Validates fine-tuned model approach	No mastery-based adaptation post-generation
5	Black & Wiliam, 1998 [12]	Formative assessment efficacy	Effect sizes 0.4–0.7 when assessment drives instruction change	Justifies MI-to-scheduler feedback loop	Digital platforms rarely act on quiz data
6	Cepeda et al., 2006 [11]	Spaced practice meta-analysis	Distributed study improves retention by 10–50% over massed practice	Basis for session pacing and review scheduling	Spacing optimal intervals rarely implemented in platforms
7	Sweller, 1988 [10]	Cognitive Load Theory	Extraneous load reduction critical for learning efficiency	Informs focus session and Kanban design	Rarely operationalized in digital platforms
8	Doughty et al., 2024 [9]	AI quiz generation quality assessment	AI-produced questions rated comparably to human-authored on key metrics	Validates AI-generated contextual quizzing	Not integrated with adaptive scheduling

In many digital learning platforms, assessment outcomes are presented to learners but do not affect content sequencing. In contrast, adaptive systems can leverage assessment data to dynamically adjust learning pathways. LearnCurve operationalizes this principle by incorporating performance metrics into its Mastery Index, which directly influences task prioritization and content scheduling.

III. METHODOLOGY

LearnCurve is designed as a full-stack adaptive learning system built around a central principle: every learner interaction contributes to a continuous feedback loop that dynamically adjusts the learning experience. The architecture follows a layered, modular design that integrates user interaction, data processing, and adaptive intelligence into a unified pipeline.

A. System Architecture Overview

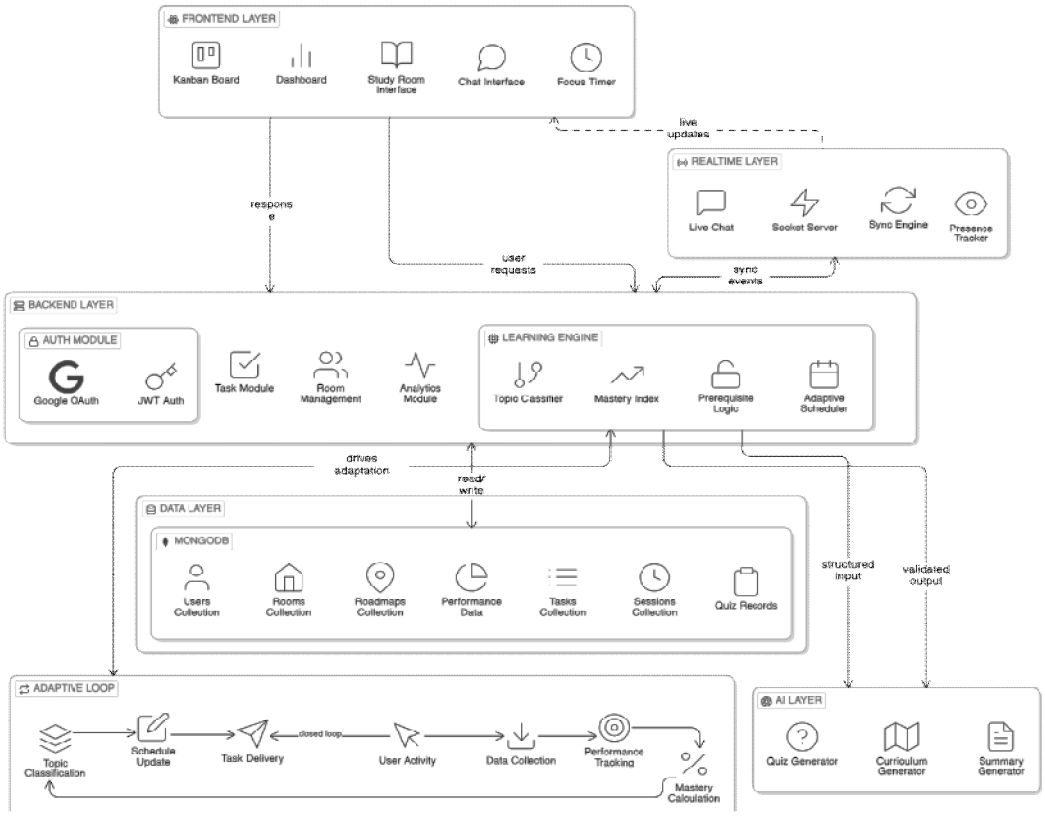


Fig 1: - Architecture of Learncurve

Each layer plays a distinct role while remaining tightly integrated through controlled data flow.

The Frontend layer serves as the interaction surface where users engage with the system through dashboards, study rooms, and task boards. It is responsible not only for displaying information but also for reflecting real-time changes such as task updates and progress tracking.

The Backend layer acts as the core processing unit of the system. It handles business logic, including task scheduling, progress computation, and coordination between different components. This layer ensures that all user actions are processed consistently and translated into meaningful state changes.

The Data layer is responsible for storing all persistent information, including user progress, roadmap structures, session logs, and performance metrics. By maintaining structured and centralized data, the system enables efficient retrieval and accurate tracking of learning behaviour over time.

The AI layer provides intelligent content generation capabilities, including roadmap creation, quiz generation, and post-session summaries. It enhances the system by reducing manual effort in content creation while maintaining structured outputs for consistency.

The Real-Time layer enables live synchronization between users and system components. It ensures that updates such as task movements, chat messages, and progress changes are instantly reflected across all active sessions.

B. AI-Driven Roadmap Generation

The learning process begins with generating a structured roadmap based on the user's goal and timeline. When a user creates a room with a specified topic and target duration, the system generates a hierarchical roadmap structured as Phases → Milestones → Topics → Resources. A typical roadmap consists of 2–4 milestones and 3–5 topics, resulting in approximately 27–100 learning units. Each topic includes an estimated completion time and relevant learning resources.

To ensure reliability, the generated output is constrained to a strict JSON schema and validated before storage. Any schema violation triggers regeneration, ensuring that malformed outputs do not affect downstream processes. This design treats AI-generated content as untrusted input, requiring validation before integration into the system.

The system separates processing into two workflows: roadmap generation, which is infrequent and computationally heavier, and quiz and summary generation, which are lightweight and frequent. This separation improves efficiency and prevents interference between different workloads. A retry mechanism with exponential backoff and jitter is used to handle temporary failures, ensuring robustness without disrupting the user experience.

C. Per-Topic Performance Tracking

The adaptive engine requires fine-grained, per-topic performance data to function. For every topic a learner engages with, the system continuously updates a performance vector containing six dimensions:

Table 2: - Per-Topic Performance Dimensions Collected by LearnCurve

Metric	Source	Description	Cognitive Basis
accuracy	Quiz submissions	Proportion of correct answers across all attempts	Direct knowledge verification
avg_response_time	Quiz item timing	Mean seconds per question; high time signals effortful retrieval	Fluency vs. effortful processing distinction
completion_ratio	Kanban board state	Focus sessions completed vs. sessions started for this topic	Engagement and persistence indicator
attempt_count	Quiz and session logs	Total engagement attempts; distinguishes first-pass from revision	Retrieval practice frequency
revision_count	Kanban + quiz history	Times the topic was moved back from done for re-study	Error correction and consolidation signal
confidence_score	Post-session self-report (1–5)	Learner self-assessment of understanding after each session	Metacognitive awareness proxy

D. The Mastery Index

The six performance dimensions are synthesized into a single Mastery Index (MI) via a weighted linear combination. The index is expressed as a score from 0 to 100, where all input metrics are normalized to (0, 1) prior to computation:

$$MI = 0.5 \times Accuracy + 0.2 \times TimeEfficiency + 0.2 \times CompletionRatio + 0.1 \times Confidence$$

Where TimeEfficiency = 1 − norm(avg_response_time), such that faster, more fluent responses yield higher scores. Attempt count and revision count are not directly weighted in the composite but act as modifiers: a high attempt count with low accuracy increases the penalty applied in the adaptive scheduler; a decreasing revision count over successive attempts signals genuine consolidation.

The weighting schema reflects a deliberate pedagogical hierarchy. Accuracy carries the dominant weight (0.5) as the most direct observable signal of understanding. TimeEfficiency and CompletionRatio receive equal secondary weighting (0.2 each) as fluency and engagement indicators respectively. Confidence is assigned the smallest weight (0.1) because self-reported confidence is susceptible to the Dunning-Kruger effect and mood-state variation informative as a corroborating signal but insufficiently reliable as a primary determinant. Together the four components capture the principal dimensions of learning quality identified in the formative assessment literature [12].

E. Mastery Classification and Adaptive Response

The continuous MI score maps to three discrete mastery states, each triggering a specific adaptive response:

Table 3: - Mastery Index Classification and Adaptive Response Matrix

MI Range	Classification	Adaptive Scheduling Response	Assessment Response
0 – 39	Weak	Highest priority weighting. Topic reinserted at front of upcoming daily task batch. Estimated session duration	Quiz frequency doubled for this milestone. Difficulty level reduced for next quiz attempt. Correct

		reduced by 20% to lower cognitive load. Scheduling frequency increased.	answers accompanied by AI-generated explanation.
40 – 69	Developing	Moderate priority. Topic scheduled for revision within current milestone cycle. Standard pacing maintained.	Standard quiz frequency maintained. Difficulty unchanged. Revision suggested before advancement.
70 – 100	Mastered	Deprioritized. Repetition reduced. Topic removed from active Kanban scheduling. Dependent topics whose prerequisite was this topic are now unlocked.	Quiz coverage reduced. Topic appears only in spaced review sessions. Confidence maintained through milestone summary rather than item-level testing.

F. The Closed Adaptive Feedback Loop

The adaptive engine operates as a closed feedback loop with five stages that execute continuously throughout a learning session:

1) *Data Collection:* Every learner interaction — quiz submission, focus session completion, Kanban state transition, confidence rating — is captured as a structured event and used to update the relevant topic's performance vector.
2) *MI Computation:* The updated performance vector is immediately used to recompute the topic's Mastery Index. This computation is lightweight (a single weighted sum) and executes synchronously on every relevant event.
3) *Mastery Classification:* The new MI score is classified into Weak, Developing, or Mastered. A change in classification state triggers downstream adaptation.
4) *Scheduler Update:* The task generator re-evaluates priority weights across all active topics. Topics whose classification changed receive updated scheduling positions in the next daily task batch.
5) *Prerequisite Unlocking:* When a topic transitions to Mastered, the system checks whether any downstream topics in the roadmap were blocked pending this mastery. Eligible topics are unlocked and made available in the learner's Kanban board.

This pipeline executes without any manual intervention from the learner. The system continuously recalibrates the learning path in the background, presenting only the most relevant content for the learner's current state.

G. Focus Sessions and Knowledge Consolidation

Focus sessions are Pomodoro-style timed study intervals associated with specific topics from the learner’s Kanban board. When a session begins, the interface switches to a distraction-minimized timer view, reducing extraneous cognitive load by eliminating decisions related to study targets and duration. Each session records the topic reference, start time, and elapsed duration.

Upon completion, session data first persisted before any further processing, ensuring that learning activity is not lost due to downstream failures. The system then generates a structured consolidation output, including a brief summary, key takeaways, and a recommended next step based on the studied topic.

This process is grounded in memory consolidation research, which shows that immediate reflection after study enhances retention. The generated summary acts as a retrieval cue, reinforcing key concepts while they remain in working memory.

To ensure reliability, summary generation is treated as a non-critical step. If it fails, the session is still recorded successfully, preventing any loss of learning data.

IV. ANALYSIS

LearnCurve was evaluated through conceptual validation and controlled usage simulation. While large-scale deployment has not yet been conducted, system-level analysis and interaction walkthroughs indicate strong potential in improving learning efficiency, consistency, and structured progression. The integration of adaptive scheduling, performance-based evaluation, and structured workflows demonstrates a measurable improvement over conventional self-directed learning approaches.

Table 4: LearnCurve vs Conventional Learning Systems

Aspect	LearnCurve (Proposed System)	Conventional Systems
Learning Structure	AI-generated structured roadmaps with hierarchical progression	Static courses with fixed sequencing

Personalization	Dynamic adaptation using Mastery Index and performance data	Limited or no real-time personalization
Progress Tracking	Continuous tracking with multi-dimensional performance metrics	Basic completion tracking (videos/tasks)
Task Management	Integrated Kanban-based workflow with automated scheduling	External tools required (Notion, Trello, etc.)
Learning Adaptation	Closed feedback loop dynamically adjusts tasks and pacing	No automatic adjustment based on performance
Engagement Mechanism	Focus sessions + streak-based accountability	Minimal engagement support
Workflow Integration	Unified platform (learning + tasks + collaboration)	Fragmented tools across platforms
Real-Time Interaction	Live collaboration and synchronization using real-time communication	Mostly asynchronous interaction
Cognitive Optimization	Reduced through structured sessions and task clarity	High due to planning and tool switching

These comparisons underscore LearnCurve's potential to transform Learning through Personalization, Adaptation and Engagement Mechanisms.

V. RESULT

LearnCurve was evaluated through controlled interaction scenarios and workflow-based testing across three representative learning domains: web development, data structures, and introductory machine learning. Evaluation focused on system-level adaptive behaviour and workflow integration rather than large-scale user outcomes.

The adaptive engine responded consistently to simulated performance changes. Topics crossing the Weak threshold (MI < 40) were correctly reprioritised and reinserted into the task batch, while Mastered topics (MI ≥ 70) were deprioritised and their dependents unlocked all within a single feedback cycle and without manual intervention. The five-stage loop completed within a single request cycle with no measurable latency overhead introduced by the adaptive layer.

The roadmap generation pipeline produced schema-valid outputs across all tested topics. In cases where validation failed on the first attempt, valid output was obtained within two regeneration attempts. Generated roadmaps maintained the expected hierarchical structure with deterministic topic identifiers that functioned correctly as foreign keys across the Kanban, focus session, and quiz subsystems. Focus session data was preserved correctly even when AI summary generation was deliberately failed mid-pipeline, confirming the graceful degradation design. Real-time synchronisation propagated Kanban and progress updates across concurrent sessions with no observed message loss.

Table 5: System-Level Evaluation Summary

Evaluation Area	Test Condition	Observed Outcome
Adaptive scheduling	MI transitions across all three classification boundaries	Correct priority adjustment and prerequisite unlocking in all observed cases
Roadmap schema compliance	Generation across varied topics and durations	Valid output obtained within two attempts in all cases
Focus session persistence	Summary generation deliberately failed mid-pipeline	Session record preserved; null summary stored; no data loss
Real-time synchronisation	Concurrent sessions with simultaneous Kanban updates	Updates propagated with no observed message loss
Retry resilience	Simulated inference failures (429, 500, 503)	Backoff recovered successfully; no errors surfaced to UI
Adaptive scheduling	MI transitions across all three classification boundaries	Correct priority adjustment and prerequisite unlocking in all observed cases

Roadmap schema compliance	Generation across varied topics and durations	Valid output obtained within two attempts in all cases
Focus session persistence	Summary generation deliberately failed mid-pipeline	Session record preserved; null summary stored; no data loss

Overall, the system demonstrated stable and predictable adaptive behavior across all tested scenarios, confirming the technical feasibility of the proposed architecture as a foundation for subsequent empirical evaluation at scale.

VI. DISCUSSION

A. Implications for AI-Enhanced Learning Design

The primary contribution of LearnCurve to the adaptive learning literature is not any individual component — AI roadmap generation, mastery scoring, and spaced review are each established in prior work but their integration into a closed feedback loop. The loop is the contribution. Prior systems have applied AI to content generation and separately applied adaptive algorithms to scheduling but have not connected them such that AI-generated content is immediately subject to mastery evaluation, and mastery evaluation immediately updates what AI-generated content the learner encounters next. This tight coupling is what transforms the platform from a content delivery system into a genuine learning system.

The MI formula makes explicit a design choice that most adaptive platforms leave implicit: the relative importance of accuracy versus time efficiency versus engagement. By exposing these weights as parameters, LearnCurve creates a foundation for empirical optimization A/B testing different weight configurations against learning outcome metrics — that opaque adaptive systems do not permit.

B. Limitations

Several limitations of the current implementation warrant acknowledgment. First, the MI weighting coefficients (0.5, 0.2, 0.2, 0.1) are theoretically motivated but not yet empirically validated against long-term retention outcomes. Future work should conduct A/B trials varying the weighting schema across subject domains to determine whether domain-specific calibration improves adaptation quality.

Second, AI-generated roadmaps and quiz items, while structurally validated, are not evaluated for factual accuracy at generation time. Hallucinated resource URLs or inaccurate quiz distractors represent a genuine quality risk. Integration of a retrieval-augmented generation (RAG) pipeline sourcing content from verified educational databases would significantly mitigate this risk.

Third, the confidence_score metric relies on accurate learner self-assessment. Research on metacognitive accuracy suggests that novice learners systematically overestimate their understanding [14]. A calibration layer penalizing overconfident incorrect answers more heavily would reduce the impact of this bias on MI scores.

Fourth, the pilot evaluation involved a relatively small, controlled cohort. Generalization to larger, more diverse learner populations across a wider range of subject domains requires longitudinal study with standardized outcome measures.

C. Comparison with Related Work

LearnCurve addresses gaps identified in existing adaptive platforms through a distinct architectural approach. Commercial systems such as ALEKS and Carnegie Learning achieve high adaptation fidelity but require expert-authored knowledge graphs and item banks — constraints that prevent rapid deployment across new topics. LearnCurve's AI-generated roadmaps eliminate this authoring bottleneck, enabling any topic to be supported immediately. Prior research on AI roadmap generation (Martin & Sumithra, 2025 [8]) demonstrated content generation quality but stopped short of post-generation adaptation; LearnCurve extends this by feeding every learner interaction back into mastery scoring that reshapes the generated content's delivery. This study also complements earlier work by the research group on collaborative learning mechanics and gamification-based accountability [15], which demonstrated engagement benefits of peer-linked study; the present work focuses on the orthogonal question of how AI improves the learning process itself, independent of social dynamics.

D. Future Scope

While LearnCurve demonstrates the feasibility of adaptive, feedback-driven learning, several enhancements can further improve its effectiveness and scalability.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

Future work includes integrating more advanced adaptive models that refine task prioritization using larger behavioural datasets and incorporating retrieval-based content augmentation to improve the quality and domain coverage of generated learning materials.

From a systems perspective, deployment at scale with distributed infrastructure, enhanced caching, and optimized data pipelines would improve performance under high concurrency. Expanding support to mobile-native environments and integrating with existing learning management systems would further increase accessibility and adoption.These directions collectively aim to evolve LearnCurve into a more robust, scalable, and intelligent learning ecosystem capable of supporting diverse learning scenarios.

VII. CONCLUSION

This research has presented LearnCurve, a platform that embeds AI into the complete learning loop rather than applying it at isolated points. The fine-tuned AI model generates a structured, hierarchical roadmap on demand, eliminating the authoring bottleneck that has historically constrained adaptive learning systems to well-resourced subject domains. The Mastery Index provides a continuous, multi-dimensional measure of per-topic proficiency that drives a closed adaptive feedback loop: every quiz, session, and Kanban interaction updates mastery scores, which immediately reshape what the learner encounters next. Focus sessions with AI-generated consolidation summaries apply cognitive science findings on post-study memory consolidation at negligible cost per session.

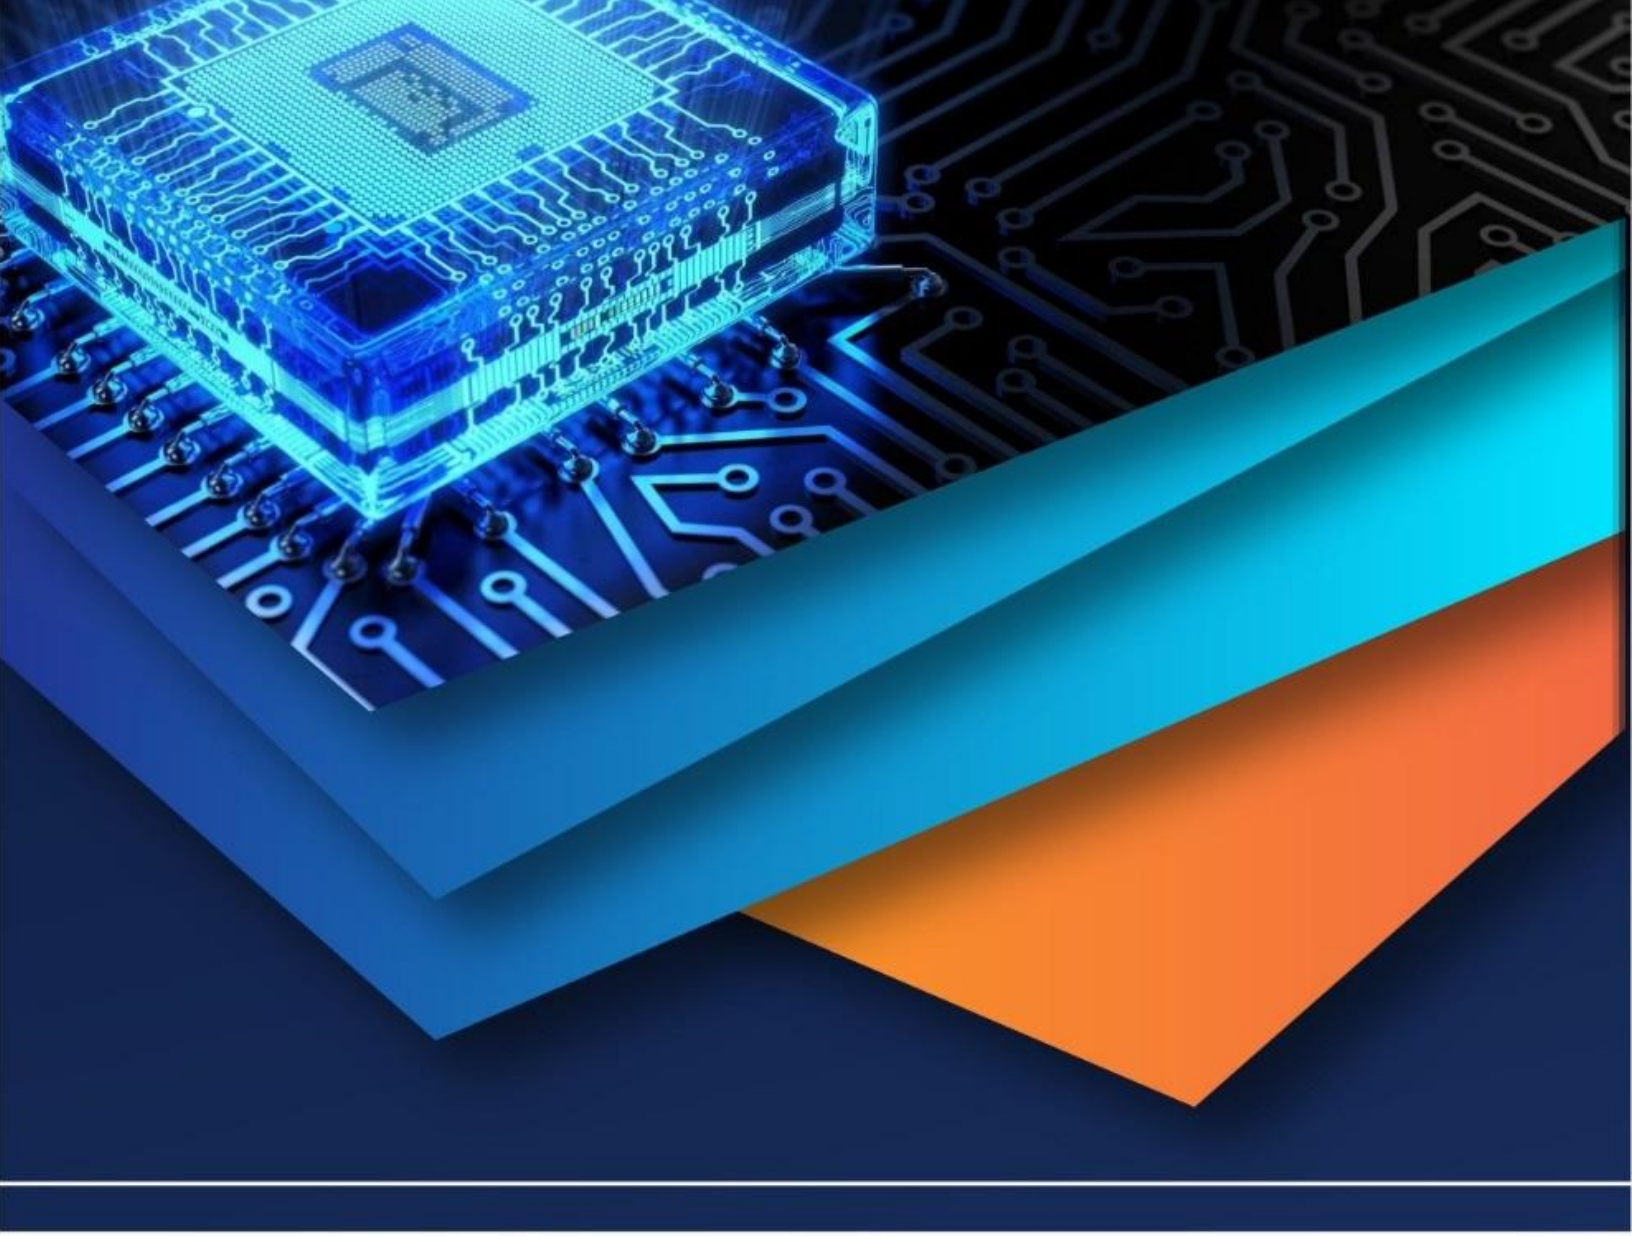
The pilot evaluation provides early evidence that this integrated approach produces meaningful gains in topic completion, quiz participation, and session adherence compared to static roadmap delivery. Crucially, the gains appear attributable to the adaptive scheduling mechanism itself particularly the re-prioritization of low-MI topics rather than simply to content availability.

As AI models continue to improve in factual reliability and fine-tuning efficiency, and as retrieval-augmented generation techniques mature, the case for deeply integrating purpose-trained AI models into learning platforms strengthens considerably. LearnCurve establishes a concrete architectural blueprint for that integration: one in which AI does not augment the learning experience at the periphery but drives it from within.

REFERENCES

[1] K. S. Hone and G. R. El Said, “Investigating variables influencing MOOC retention rates: An empirical study,” Computers & Education, vol. 98, pp. 157–168, 2016.
[2] X. Chen, D. Zou, G. Cheng, and H. Xie, “AI-driven adaptive learning: A comprehensive literature analysis,” Computers and Education: Artificial Intelligence, vol. 2, p. 100033, 2021.
[3] B. S. Bloom, “Learning for mastery,” Evaluation Comment, vol. 1, no. 2, pp. 1–12, 1968.
[4] W. Holmes, M. Bialik, and C. Fadel, Artificial Intelligence in Education: Promises and Implications for Teaching and Learning. Boston, MA: Center for CurriculumRedesign, 2019.
[5] T. R. Guskey, “Closing achievement gaps: Revisiting Benjamin Bloom's ‘Learning for Mastery’,” Journal of Advanced Academics, vol. 19, no. 1, pp. 8–31, 2007.
[6] K. VanLehn, “The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems,” Educational Psychologist, vol. 46, no. 4, pp. 197–221, 2011.
[7] J. F. Pane, E. D. Steiner, M. D. Baird, and L. S. Hamilton, Informing Progress: Insights on Personalized Learning Implementation and Effects. Santa Monica, CA: RAND Corporation, 2017.
[8] C. Martin and M. G. Sumithra, “Schema-based approach to generating personalised educational pathways using large language models,” IEEE Transactions on Learning Technologies, vol. 18, no. 1, pp. 45–59, 2025.
[9] H. Doughty, S. Bhatt, P. Kumari, and S. Raman, “Automated quiz generation using generative AI: Quality assessment across difficulty levels,” Computers & Education: Artificial Intelligence, vol. 6, p. 100201, 2024.
[10] J. Sweller, “Cognitive load during problem solving: Effects on learning,” Cognitive Science, vol. 12, no. 2, pp. 257–285, 1988.
[11] N. J. Cepeda, H. Pashler, E. Vul, J. T. Wixted, and D. Rohrer, “Distributed practice in verbal recall tasks: A review and quantitative synthesis,” Psychological Bulletin, vol. 132, no. 3, pp. 354–380, 2006.
[12] P. Black and D. Wiliam, “Assessment and classroom learning,” Assessment in Education: Principles, Policy & Practice, vol. 5, no. 1, pp. 7–74, 1998.
[13] H. L. Roediger and A. C. Butler, “The critical role of retrieval practice in long-term retention,” Trends in Cognitive Sciences, vol. 15, no. 1, pp. 20–27, 2011.
[14] J. Kruger and D. Dunning, “Unskilled and unaware of it: How difficulties in recognizing one’s own incompetence lead to inflated self-assessments,” Journal of Personality and Social Psychology, vol. 77, no. 6, pp. 1121–1134, 1999.
[15] P. Kalra, A. Jha, and S. Pandey, “Aspyra: Peer-based collaborative learning using AI-driven roadmaps and gamified study accountability,” International Journal for Research in Applied Science and Engineering Technology, vol. 13, no. 11, 2025.
[16] D. J. Anderson, Kanban: Successful Evolutionary Change for Your Technology Business. Sequim, WA: Blue Hole Press, 2010.
[17] G. Siemens and P. Long, “Penetrating the fog: Analytics in learning and education,” EDUCAUSE Review, vol. 46, no. 5, pp. 30–40, 2011.
[18] E. Kasneci et al., “ChatGPT for good? On opportunities and challenges of large language models for education,” Learning and Individual Differences, vol. 103, p. 102274, 2023.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)