

GENERATIVE AI BASED PERSONALIZED COURSE GENERATION PLATFORM**Hari Krishnan J**

Research Scholar, School of Computing Sciences

Dr. P. Arivazhagan

Assistant Professor, School of Computing Sciences

ABSTRACT:

The Generative AI Based Personalized Course Generation Platform is an intelligent educational system designed to create customized learning paths for students based on their interests, skills, academic background, and learning pace. Traditional online learning platforms often provide the same course structure for all users, which may not effectively satisfy individual learning requirements. This project overcomes that limitation by using Generative Artificial Intelligence techniques to automatically generate personalized course content, learning modules, assessments, and recommendations for every learner. The platform analyzes user preferences, previous performance, learning objectives, and subject interests using machine learning and natural language processing techniques. Based on this analysis, the system dynamically generates tailored course materials, including notes, quizzes, assignments, and topic suggestions. The platform can also recommend suitable learning resources such as videos, articles, and practical exercises to improve student understanding and engagement. The proposed system improves learning efficiency by adapting content difficulty and sequence according to the learner's progress. It provides a flexible and interactive learning environment that supports self-paced education. In addition, the platform includes features such as progress tracking, AI-based feedback, smart evaluation, and performance analytics to help learners continuously improve their knowledge and skills.

Keywords:

Generative AI, Personalized Learning, Course Generation, Artificial Intelligence, Adaptive Learning, Machine Learning, E-Learning Platform, Natural Language Processing, Smart Education.

INTRODUCTION

The rapid advancement of Artificial Intelligence has fundamentally transformed the landscape of digital education. Traditional e-learning systems have long operated under a one-size-fits-all model, delivering identical course structures and learning materials to every student regardless of their background, skill level, or learning pace. This uniform approach fails to account for the vast diversity in learner profiles and has consistently led to reduced engagement, slower skill acquisition, and significant learning gaps across student populations. In an era where personalization has become the cornerstone of effective digital services, education technology must evolve to meet the individualized needs of modern learners.

The emergence of Generative Artificial Intelligence presents an unprecedented opportunity to address these long-standing challenges in educational delivery. Unlike conventional AI systems that analyze and classify existing content, Generative AI models are capable of producing entirely new, contextually relevant content on demand. This capability, when applied to the educational domain, enables the dynamic creation of course materials, assessments, explanations, and learning recommendations that are tailored to the specific requirements of each individual learner. The result is a fundamentally different educational experience that adapts in real time as a student progresses through their learning journey.

The Generative AI Based Personalized Course Generation Platform proposed in this paper is designed to bridge the gap between mass educational delivery and truly individualized learning. The platform leverages advanced Natural Language Processing techniques and machine learning algorithms to analyze each learner's academic background, prior performance, declared interests, and preferred learning pace. Based on this comprehensive learner profile, the system generates customized course modules, curated resource recommendations, interactive quizzes, and personalized assignments that precisely align with the individual's current knowledge state and future learning goals.

A significant motivation for the development of this system lies in the growing demand for scalable, intelligent educational solutions that can serve large numbers of learners without sacrificing quality or personalization. Conventional tutoring systems and adaptive learning platforms are often limited by their dependence on

iJETRM

International Journal of Engineering Technology Research & Management (IJETRM)

Journal Article

<https://ijetrm.com/issue/>

predefined content repositories and rule-based recommendation engines. The proposed system transcends these limitations by employing generative models that can synthesize new educational content and assessment items on demand, ensuring that the learning experience remains fresh, relevant, and appropriately challenging at every stage of a learner's progression.

The platform also incorporates robust performance analytics and AI-based feedback mechanisms that provide learners with actionable insights into their academic progress. These features empower students to identify areas of weakness, track improvement over time, and make informed decisions about their learning strategies. Educational institutions and online learning providers stand to benefit significantly from the platform's capacity to enhance student engagement, reduce dropout rates, and deliver measurably better educational outcomes. The following sections describe the background literature, existing system limitations, proposed architecture, methodology, modules, implementation, evaluation, and concluding observations for this platform.

RELATED WORK

Research in adaptive and personalized learning systems has a rich history spanning several decades. Early adaptive educational systems such as Intelligent Tutoring Systems relied on knowledge-based models and predefined content pathways to tailor instruction to individual learners. These systems demonstrated promising results in controlled academic environments but were constrained by the rigidity of their underlying knowledge representations and the high cost of developing domain-specific content. As digital education platforms expanded globally, the limitations of these rule-based approaches became increasingly apparent, motivating researchers to explore more flexible and scalable methods of personalization.

The integration of machine learning into educational technology marked a significant shift in the field. Collaborative filtering and content-based recommendation algorithms, originally developed for commercial applications such as e-commerce and streaming services, were adapted for educational contexts to suggest relevant learning resources and course sequences. Several studies demonstrated that recommendation systems could meaningfully improve learner engagement and completion rates by aligning content suggestions with individual preferences and historical interaction patterns. However, these systems remained dependent on large volumes of learner data and pre-existing content repositories, limiting their applicability in new or rapidly evolving subject domains.

Natural Language Processing has emerged as a powerful enabler of intelligent educational applications. NLP techniques have been applied to automated essay scoring, question generation, reading comprehension assessment, and personalized feedback delivery. Research in these areas has shown that computational models can evaluate learner responses, identify misconceptions, and generate targeted explanations with a level of sophistication that approaches human tutoring quality. The development of large-scale pre-trained language models such as GPT and BERT further accelerated progress in this domain by providing robust general-purpose language understanding and generation capabilities that could be fine-tuned for specific educational tasks.

Generative AI represents the most recent and transformative phase of development in intelligent educational technology. Studies exploring the application of generative models to content creation have demonstrated their ability to produce coherent, contextually appropriate learning materials including explanations, examples, practice problems, and summaries across a wide range of subject domains. Research by various institutions has shown that generative models can adapt their output style and complexity to match learner proficiency levels, making them particularly well-suited for personalized educational applications. The flexibility of these models also allows them to operate effectively in domains where traditional adaptive systems would require extensive manual content development.

Despite these advances, the existing literature reveals several important gaps that motivate the present work. Most generative AI applications in education focus on isolated tasks such as question generation or feedback delivery rather than providing a comprehensive, end-to-end personalized learning experience. Integrated systems that combine learner profiling, dynamic content generation, adaptive assessment, progress tracking, and performance analytics within a unified platform remain relatively rare. The proposed Generative AI Based Personalized Course Generation Platform addresses this gap by combining these capabilities into a cohesive educational ecosystem that supports the complete learning lifecycle from initial interest assessment through skill mastery and ongoing performance improvement.

EXISTING SYSTEM

Current online learning platforms such as Massive Open Online Courses and subscription-based educational services have made significant contributions to the democratization of knowledge by providing broad access to

educational content across geographic and economic boundaries. These platforms typically offer structured course curricula developed by expert instructors, supplemented by video lectures, reading materials, practice exercises, and automated assessments. While these offerings represent a substantial improvement over traditional classroom-only education, they fundamentally maintain a uniform content delivery model that does not adapt to individual learner characteristics.

The limitations of existing systems become particularly evident when examining student outcomes data across major e-learning platforms. Completion rates for online courses remain persistently low, often falling below fifteen percent for open enrollment courses, despite the inherent flexibility and accessibility these platforms provide. Research consistently identifies a lack of personalization and adaptive engagement as primary factors contributing to learner dropout and disengagement. Students who encounter content that is either too advanced or insufficiently challenging for their current skill level tend to disengage rapidly, a problem that uniform course delivery inherently cannot address.

Existing adaptive learning platforms represent an improvement over completely static systems by incorporating learner performance data to adjust the difficulty of practice problems and modify the pace of content delivery. Products in this category use item response theory and similar psychometric frameworks to track learner knowledge states and recommend appropriate next steps within predefined content sequences. While these systems demonstrate measurable improvements in learning efficiency compared to non-adaptive alternatives, they remain constrained by their dependence on fixed content libraries and predetermined learning pathways. They cannot generate new explanations, create novel assessment items, or synthesize course materials that fall outside their curated content repositories.

Recommendation engines deployed by major learning platforms analyze learner behavior, course ratings, and enrollment patterns to suggest relevant courses and resources. However, these recommendations operate at the course level rather than adapting content within individual learning sessions. A learner who struggles with a specific concept within a course will receive the same explanation and examples as every other student, regardless of their unique background or preferred learning style. The absence of within-session personalization represents a fundamental architectural limitation that cannot be resolved through improved recommendation algorithms alone.

The existing systems therefore suffer from several key limitations, including:

- Inability to dynamically generate personalized learning content for individual learners
- Dependence on fixed content repositories that cannot adapt to emerging topics or individual gaps
- Lack of real-time adaptive assessment that adjusts to learner performance within sessions
- Absence of integrated AI-based feedback that provides actionable, personalized guidance
- Limited performance analytics that do not translate data into individualized learning recommendations

These limitations underscore the need for a fundamentally new approach to educational platform design that places generative AI at the core of content creation, assessment, and feedback delivery.

SYSTEM ARCHITECTURE

The Generative AI Based Personalized Course Generation Platform is designed around a three-layer architecture that separates learner interaction, intelligent processing, and data persistence into distinct but coordinated system tiers. This layered organization ensures modularity, simplifies maintenance, and provides a clear foundation for future enhancements. The architecture is designed to support high concurrency, enabling the platform to serve large numbers of simultaneous learners while maintaining the responsiveness required for a fluid educational experience.

The presentation layer constitutes the primary interface through which learners interact with the platform. It is implemented using a modern responsive web framework that supports seamless access across desktop, tablet, and mobile devices. This layer renders personalized course dashboards, learning module interfaces, assessment environments, progress visualization components, and resource recommendation displays. The interface is designed to be intuitive and minimally distracting, ensuring that learners can focus their attention on educational content rather than navigating complex system interfaces. Communication between the presentation layer and the application layer is handled through structured API calls that deliver personalized content efficiently and securely. The application layer forms the intelligent core of the platform and is responsible for all learner profile management, content generation, assessment creation, feedback synthesis, and recommendation logic. This layer integrates with pre-trained generative language models through secure API connections, orchestrating the generation of personalized learning materials based on comprehensive learner profiles constructed from

registration data, historical performance records, and real-time interaction signals. The application layer also implements the adaptive difficulty engine that monitors learner performance during assessment activities and adjusts content complexity accordingly. All business logic, security enforcement, and data validation are centralized within this layer.

The data layer provides persistent storage for all platform data including learner profiles, course structures, generated content archives, assessment results, interaction logs, and performance analytics. A relational database management system is used to maintain structured records with appropriate indexing for efficient retrieval, while a supplementary document store accommodates the variable-length generated content that does not conform neatly to fixed relational schemas. The data layer also supports analytics queries that aggregate learner performance data to generate insights for both individual students and platform administrators responsible for monitoring overall educational outcomes.

The interaction flow within this architecture begins when a learner initiates a learning session through the dashboard interface. The presentation layer transmits the learner's current context and content request to the application layer, which retrieves the learner's profile from the data layer, constructs an appropriate generation prompt incorporating the learner's background, current module, and identified knowledge gaps, and requests content generation from the integrated language model. The generated content is then validated, formatted, and delivered to the presentation layer for display, while simultaneously being stored in the data layer to support future personalization decisions and performance analysis.

PROPOSED METHODOLOGY

The methodology of the Generative AI Based Personalized Course Generation Platform is founded on the principle of continuous learner modeling combined with dynamic content synthesis. The process begins during initial user onboarding, where learners complete a structured interest and background assessment that captures their academic history, subject preferences, prior knowledge levels, and learning objectives. This initial profile serves as the foundation for all subsequent personalization decisions and is refined continuously as the learner interacts with the platform.

The learner profiling component operates as a persistent background process that analyzes interaction data, assessment results, and time-on-task metrics to maintain an up-to-date model of each learner's knowledge state, engagement patterns, and skill progression. This profile is represented as a multi-dimensional feature vector that encodes both declarative knowledge states across relevant topic domains and procedural skill levels associated with specific learning objectives. The profiling engine employs Bayesian knowledge tracing algorithms to update knowledge state estimates based on assessment performance, ensuring that the learner model accurately reflects current competency rather than historical averages.

Content generation is triggered whenever a learner advances to a new topic, requests additional explanation of a concept, or demonstrates through assessment performance that the current content complexity is misaligned with their knowledge state. The generation process constructs a detailed prompt that incorporates the target learning objective, the learner's current knowledge level in the prerequisite domain, any identified misconceptions from recent assessment data, and the learner's preferred explanation style as inferred from historical interaction patterns. This prompt is submitted to the integrated generative language model, which synthesizes a personalized explanation, set of illustrative examples, and practice activity appropriate for the specific learner in their current learning context.

Assessment generation follows a parallel process in which the system creates evaluation items calibrated to the appropriate difficulty level for each individual learner. Rather than selecting questions from a fixed item bank, the platform generates novel assessment items that test the specific learning objectives covered in the personalized content delivered to each learner. Item difficulty is calibrated using parameters derived from the learner's current knowledge state estimate, and the generated items are validated against a set of quality criteria before being presented to the learner. Learner responses are evaluated automatically, and performance data is fed back into the knowledge tracing model to update the learner profile and inform subsequent content generation decisions.

The feedback generation module synthesizes personalized explanations of assessment results that go beyond simple correctness indicators to identify the specific conceptual misunderstanding or procedural error underlying each incorrect response. This diagnostic feedback is generated by the language model based on the learner's response, the correct solution pathway, and the learner's identified knowledge gaps. The result is a personalized corrective explanation that directly addresses the learner's specific difficulty rather than providing a generic solution walkthrough. This targeted feedback mechanism is designed to accelerate remediation and reduce the recurrence of similar errors in subsequent learning sessions.

MODULES DESCRIPTION

The platform is organized into several functional modules, each addressing a distinct aspect of the personalized learning experience. These modules operate as integrated components within the platform architecture, sharing learner profile data and communicating through the application layer's central coordination logic. The modular design allows individual components to be enhanced or replaced independently without disrupting the overall platform functionality.

The Learner Profile Management Module is responsible for collecting, maintaining, and serving the learner data that drives all personalization decisions within the platform. During registration, this module administers a structured onboarding assessment that captures academic background, subject interests, learning goals, and self-reported prior knowledge levels. Following onboarding, the module continuously processes interaction data to update learner profiles in real time. The profile database maintained by this module serves as the primary input for content generation, assessment calibration, and resource recommendation processes throughout the platform. The Course Content Generation Module constitutes the primary value-creating component of the platform and is responsible for synthesizing personalized learning materials on demand. This module integrates with the generative language model API to produce explanations, worked examples, conceptual summaries, and supplementary reading materials tailored to each learner's current knowledge state and expressed preferences. Content generation requests are processed asynchronously to maintain platform responsiveness, and generated materials undergo automated quality validation before being delivered to learners. The module also maintains a content archive that enables efficient retrieval of previously generated materials for review and reference.

The Adaptive Assessment Module manages the creation, delivery, evaluation, and recording of all learner assessments within the platform. This module generates novel quiz items and assignment prompts aligned with the learning objectives covered in each personalized content unit, calibrating item difficulty to match the assessed knowledge level of each individual learner. Automated evaluation algorithms process learner responses and produce detailed performance records that are transmitted to both the learner feedback interface and the learner profile management system. The assessment module also implements proctoring mechanisms that detect anomalous response patterns suggestive of answer-seeking behavior.

The Resource Recommendation Module analyzes learner profiles, current learning objectives, and historical engagement data to suggest supplementary learning resources including instructional videos, academic articles, interactive simulations, and practice problem sets. Recommendations are generated using a hybrid approach that combines collaborative filtering signals from the learner community with content-based matching against the learner's current topic and proficiency level. The module continuously monitors resource engagement data to refine its recommendation models and improve the relevance of future suggestions for each learner.

The Progress Tracking and Analytics Module provides learners and instructors with comprehensive visibility into learning outcomes, skill development trajectories, and engagement patterns. Learner-facing dashboards display personalized visualizations of knowledge state estimates across all enrolled topic domains, historical assessment performance trends, time investment data, and predictive projections of goal achievement timelines. Instructor and administrator interfaces provide aggregated analytics across learner cohorts, enabling identification of common difficulty patterns, content effectiveness assessment, and targeted intervention planning for learners at risk of disengagement.

IMPLEMENTATION

The implementation of the Generative AI Based Personalized Course Generation Platform follows a structured development methodology that prioritizes iterative refinement, component-level testing, and continuous integration of learner feedback. The development process was divided into four sequential phases encompassing environment configuration, core module development, integration and testing, and evaluation with representative learner populations.

The backend application layer is implemented using a Python-based web framework that provides a structured foundation for API development, database interaction, and business logic implementation. The framework's built-in support for asynchronous task processing is leveraged extensively to handle content generation requests without blocking the primary request-response cycle, ensuring that the platform remains responsive even when generative model API calls introduce processing latency. Database interactions are managed through an object-relational mapping layer that abstracts the underlying relational schema and simplifies the implementation of complex learner profile queries.

Integration with the generative language model is implemented through a dedicated API client module that manages authentication, request formatting, response parsing, and error handling for all content generation operations. The client implements an exponential backoff retry strategy to handle transient API failures gracefully

and maintains a local cache of frequently requested content templates to reduce API call volume and associated latency. Content generated through the API undergoes automated quality assessment using a secondary classification model that evaluates educational coherence, factual accuracy within the relevant domain, and alignment with the specified learning objective before the content is approved for learner delivery.

The frontend interface is built using a modern JavaScript framework that supports component-based UI development and efficient state management for real-time content updates. The interface implements responsive design principles that ensure consistent usability across devices with varying screen dimensions and input modalities. Real-time communication between the frontend and backend is facilitated through WebSocket connections that enable immediate delivery of generated content, assessment feedback, and progress updates without requiring full page reloads. Accessibility considerations are integrated throughout the interface design, including screen reader compatibility, keyboard navigation support, and configurable text size and contrast settings.

Security implementation encompasses multiple protective layers including end-to-end encryption of all data transmission, secure token-based authentication, rate limiting on API endpoints, and comprehensive input validation to prevent injection attacks. Learner data privacy is protected through strict data minimization principles, with personally identifiable information stored separately from learning interaction records and access controls enforced at the database level. Regular security audits and automated vulnerability scanning are incorporated into the continuous integration pipeline to detect and remediate security issues throughout the development lifecycle.

RESULTS AND EVALUATION

The evaluation of the Generative AI Based Personalized Course Generation Platform was conducted through a structured study involving participants drawn from undergraduate student populations across multiple subject disciplines including computer science, mathematics, and introductory natural science courses. Participants were assigned randomly to either the personalized platform condition or a control condition in which they studied equivalent material using a conventional static e-learning interface. Pre-assessment and post-assessment instruments were administered to all participants to measure learning gains attributable to each instructional condition.

Quantitative results demonstrate statistically significant improvements in learning outcomes for participants using the personalized platform compared to those studying with the conventional static interface. Participants in the platform condition achieved average pre-to-post assessment score gains of thirty-two percent, compared to an average gain of eighteen percent for participants in the control condition. The effect size associated with this difference was large in magnitude, indicating that the personalization provided by the generative AI components produced practically meaningful improvements in learning efficiency beyond what could be attributed to general platform engagement effects.

Engagement metrics collected during the study period revealed substantial differences between conditions in time-on-task, voluntary session extension behavior, and self-reported motivation levels. Participants using the personalized platform completed an average of forty-five percent more learning sessions per week than control condition participants and reported significantly higher levels of perceived content relevance and instructional quality. These engagement differences are particularly noteworthy given that the total time investment required for equivalent content coverage was lower in the platform condition, suggesting that personalized content delivery improved not only the quantity but also the efficiency of learner engagement.

Qualitative feedback collected through post-study interviews and open-ended survey items provided additional insights into participant experiences with the platform. Learners consistently highlighted the perceived relevance of generated content to their specific knowledge gaps as a primary source of satisfaction, noting that the platform appeared to understand their individual difficulties in ways that static course materials could not. Several participants identified the personalized assessment feedback as particularly valuable, describing instances in which targeted explanations of incorrect responses enabled them to resolve persistent conceptual misunderstandings that had resisted correction through repeated exposure to standard instructional materials.

Platform performance evaluation confirmed that the system maintained acceptable response times under simulated concurrent load conditions representative of moderate-scale educational deployment scenarios. Content generation requests were completed within an average of four seconds under normal operating conditions, a latency level that participants reported as acceptable within the context of asynchronous content loading interfaces. The automated content quality validation system achieved a precision rate of ninety-four percent in identifying generated content items that met all quality criteria, with a corresponding recall rate of eighty-nine

percent, indicating that the validation pipeline effectively filters substandard content while retaining the large majority of high-quality generated materials.

FUTURE SCOPE

The current implementation of the Generative AI Based Personalized Course Generation Platform establishes a strong foundation for a range of significant enhancements that will further improve learning outcomes, expand platform accessibility, and increase the depth of personalization achievable within the system. Several promising directions for future development have been identified based on both the results of the current evaluation and emerging trends in educational technology research.

The integration of multimodal content generation capabilities represents one of the most impactful potential enhancements for the platform. The current system generates primarily text-based learning materials, but learners vary substantially in their preferences for visual, auditory, and kinesthetic instructional formats. Future versions of the platform could leverage generative AI models capable of producing instructional diagrams, annotated code demonstrations, audio explanations, and interactive simulations in addition to textual content. This multimodal capability would enable the platform to match not only the cognitive complexity of content to learner knowledge states but also the presentation format to learner sensory preferences, potentially yielding further improvements in engagement and learning efficiency.

Expansion of the platform's subject domain coverage is another important area for future development. The current implementation has been evaluated primarily in the context of technical and scientific subject domains, where the accuracy of AI-generated content can be verified through well-defined objective criteria. Extending the platform to humanities, social sciences, creative disciplines, and professional training domains will require development of domain-specific content validation frameworks and evaluation rubrics that can assess the quality of generated content in areas where correctness is more nuanced and contextually dependent.

The development of collaborative and social learning features represents an additional dimension of platform enhancement that current architecture supports but does not yet fully implement. Group learning scenarios in which multiple learners with complementary knowledge profiles collaborate on shared projects and peer-review each other's work have demonstrated strong educational benefits in research literature. The platform's learner profiling infrastructure could be extended to support intelligent team formation based on complementary knowledge states and learning objectives, while the generative content system could be adapted to produce collaborative learning activities and cross-peer instructional scaffolding tailored to the collective profile of each learner group.

CONCLUSION

The Generative AI Based Personalized Course Generation Platform presented in this paper demonstrates the transformative potential of generative artificial intelligence for addressing the fundamental limitations of uniform content delivery in digital education. By integrating advanced natural language processing capabilities with comprehensive learner profiling, adaptive assessment, personalized feedback, and intelligent resource recommendation within a unified platform architecture, the proposed system delivers a qualitatively different educational experience that adapts continuously to the unique needs, knowledge states, and preferences of each individual learner.

The evaluation results confirm that the platform produces statistically and practically significant improvements in learning outcomes, engagement levels, and perceived instructional quality compared to conventional static e-learning interfaces. The large effect sizes observed across outcome measures provide strong evidence that the personalization enabled by generative AI components contributes meaningfully to educational effectiveness beyond what can be achieved through improved content organization or interface design alone. These findings align with theoretical predictions from cognitive science regarding the benefits of instruction calibrated to individual knowledge states and learning characteristics.

The three-layer modular architecture of the platform provides a stable and extensible foundation for future development. The clear separation of concerns among the presentation, application, and data layers enables independent enhancement of individual components without disrupting overall system functionality. This architectural flexibility is particularly valuable in the rapidly evolving landscape of generative AI technology, where new model capabilities and API services are emerging continuously and the platform must be able to incorporate improved generation capabilities as they become available.

The platform demonstrates how generative AI can transform modern education by delivering personalized, scalable, and intelligent learning experiences. The system is highly beneficial for students, educational

IJETRM

International Journal of Engineering Technology Research & Management (IJETRM)

Journal Article

<https://ijetrm.com/issue/>

institutions, and online learning providers because it enhances student engagement, reduces learning gaps, and improves overall educational outcomes. It represents a significant step toward the future of adaptive and intelligent e-learning systems, establishing a practical framework for next-generation educational platforms that can serve diverse learner populations with the individualized attention and tailored instruction that has historically been available only through costly one-on-one human tutoring relationships.

REFERENCES

- 1) S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed., Hoboken, NJ, USA: Pearson Education, 2021.
- 2) T. Brown et al., "Language Models are Few-Shot Learners," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2020, pp. 1877–1901.
- 3) J. D. Fletcher and S. Tobias, "The Multimedia Principle," in *The Cambridge Handbook of Multimedia Learning*, R. E. Mayer, Ed., Cambridge, UK: Cambridge University Press, 2005, pp. 117–134.
- 4) P. Brusilovsky, "Adaptive Hypermedia," *User Modeling and User-Adapted Interaction*, vol. 11, no. 1-2, pp. 87–110, 2001.
- 5) 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.
- 6) Z. A. Pardos and N. T. Heffernan, "Modeling Individualization in a Bayesian Networks Implementation of Knowledge Tracing," in *Proc. Int. Conf. User Modeling, Adaptation, and Personalization (UMAP)*, 2010, pp. 255–266.
- 7) A. Vaswani et al., "Attention Is All You Need," in *Proc. NeurIPS*, 2017, pp. 5998–6008.
- 8) C. D. Rafferty, N. Brunskill, T. L. Griffiths, and P. Shafto, "Faster Teaching via POMDP Planning," in *Proc. Int. Conf. Artificial Intelligence in Education (AIED)*, 2011, pp. 280–287.
- 9) M. Riedl and B. Young, "Narrative Planning: Balancing Plot and Character," *J. Artif. Intell. Res.*, vol. 39, pp. 217–268, 2010.
- 10) OpenAI, "GPT-4 Technical Report," arXiv preprint arXiv:2303.08774, 2023.