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

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 6, June 2026

DOI: 10.17148/IJARCCE.2026.156109

Agentic AI for Personalized Learning

Afnan Abdul Raheem Nasir[1], Dr. A. Gayathri[2]

M. Tech Student, Department of Computer Science, UCEOU, Hyderabad, India[1]

Assistant Professor, Department of Computer Science, UCEOU, Hyderabad, India[2]

Abstract: Adaptive tutoring systems aim to personalize instruction by modeling a learner's current mastery and dynamically selecting the next concept to teach. This paper presents an Agentic AI framework for personalized learning that combines transformer-based knowledge tracing, autonomous decision-making, and dynamic content generation. The proposed system estimates student mastery using an attention-based knowledge tracing model, selects the next skill with a Deep Q-Network (DQN), and generates lesson and quiz content for targeted review. The architecture is modular, practical, and suitable for deployment in a web-based tutoring environment. The experimental use of student interaction data and reinforcement-based skill selection demonstrate how adaptive teaching can improve both personalization and learner engagement.

Keywords: Adaptive tutoring, knowledge tracing, reinforcement learning, DQN, AI tutor, personalized education.

I. INTRODUCTION

Personalized education is increasingly important because learners differ in prior knowledge, pace, and learning style. Traditional education systems often follow fixed curricula, which may not match individual needs. As a result, students may spend too much time on already mastered skills or struggle with concepts that require more guidance. This paper proposes an adaptive AI tutor that integrates knowledge tracing, deep reinforcement learning, and adaptive content generation. The system estimates a learner's mastery over time, decides which concept to teach next, and delivers personalized lesson content. The goal is to improve the quality of teaching by adapting the learning path to each student's current knowledge state.

A. Motivation

Many intelligent tutoring systems still rely on static or rule-based lesson sequencing. Such approaches are limited be cause they cannot fully capture changes in a learner's mastery during a session. A model that continuously updates its understanding of the learner and suggests the next best concept can provide a more effective learning experience.

B. Problem Statement

A personalized tutor should answer at least three questions continuously:

- What does the learner currently know? ;
- Which concept should be taught next? ;
- How can the system reinforce weak concepts?

The proposed system addresses these challenges through an integrated pipeline of data preprocessing, knowledge tracing, mastery estimation, DQN-based planning, and adaptive content generation.

C. Contributions

The main contributions of this paper are:

- A modular adaptive tutoring framework that combines knowledge tracing and reinforcement learning;
- A DQN-based policy for selecting the next learner skill;
- Adaptive generation of lesson and review content;
- A practical web-based tutor implementation for real-time interaction.

II. RELATED WORK

Research in educational technology has evolved from simple rule-based tutors to data-driven and machine-learning-based systems. In particular, knowledge tracing and reinforcement learning have gained attention for their ability to model learner behavior and optimize instructional decisions.

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

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁑ Peer-reviewed & Refereed journal ⁑ Vol. 15, Issue 6, June 2026

DOI: 10.17148/IJARCCE.2026.156109

A. Knowledge Tracing

Early intelligent tutoring systems used Bayesian models to estimate student mastery. Later, deep learning approaches improved this by modeling sequential interaction history. Deep Knowledge Tracing (DKT) demonstrated that recurrent neural models can capture temporal learning patterns [1], [2]. More recent work, such as Context-Aware Attentive Knowledge Tracing (AKT), improves prediction by using attention mechanisms to focus on relevant past interactions [3], [4]. The self-attentive approach further shows the benefit of modeling long-range dependencies in student learning sequences [5].

B. Reinforcement Learning in Education

Reinforcement learning has been widely used to optimize sequential decision making. The introduction of deep Q learning transformed the field by enabling agents to learn directly from high-dimensional state representations [6]. Sutton and Barto provide the standard formalization of reinforcement learning and its core algorithms [7]. In the educational domain, RL has been applied to personalized learning path recommendation and adaptive exercise selection [8], [9], [10], [11].

C. Adaptive Content Generation

Recent advances in language models have encouraged the use of AI-generated explanations, hints, and quizzes. Few-shot prompting and large language models can generate instructional content with minimal manual effort [12]. In intelligent tutoring, such methods can be combined with learner models to improve the relevance of generated learning materials [13], [14].

D. Research Gap

Although prior systems have explored knowledge tracing, reinforcement learning, and content generation separately, fewer systems combine all three into one practical adaptive tutoring architecture. This paper addresses that gap by integrating interpretable mastery estimation, DQN-based skill selection, and content generation into a unified system.

III. METHODOLOGY

The proposed system follows a pipeline that transforms raw student interactions into adaptive lesson decisions and personalized learning content.

A. Overall Workflow

The workflow consists of five stages:

1. Preprocess raw student interaction data;
2. Estimate mastery using an AKT model;
3. Represent the learner state as a mastery vector;
4. Use DQN to select the next skill;
5. Generate lesson or review content tailored to the learner.

This process repeats throughout the session, making the tutor adaptive to recent performance. A simplified view of the workflow is shown in Fig. 1.

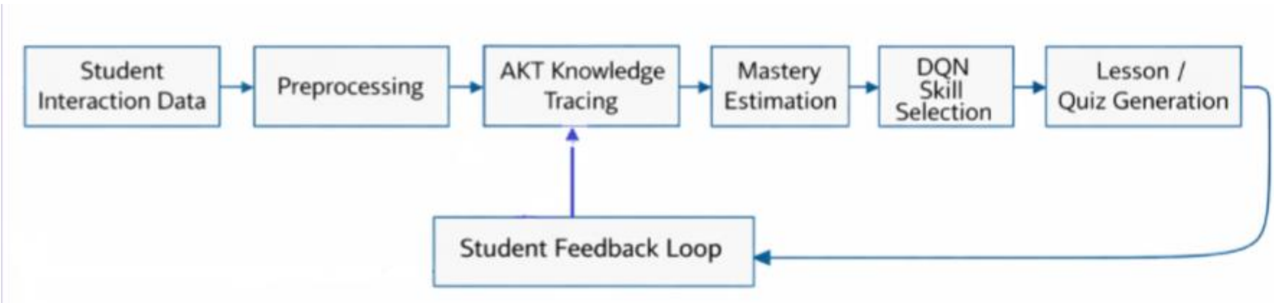

Figure 1: Overall workflow of the proposed adaptive AI tutor system.

B. Data Preparation

The system uses interaction logs containing student identifiers, skill identifiers, and correctness labels. Data preprocessing includes:

- Loading the raw CSV dataset;
- Mapping skill IDs to continuous integer indices;

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁙ Peer-reviewed & Refereed journal ⁙ Vol. 15, Issue 6, June 2026
DOI: 10.17148/IJARCCE.2026.156109

- Grouping interactions by student;
- Constructing ordered learning sequences;
- And splitting the data into training and evaluation sets.

This preprocessing enables the knowledge tracer and learning policy to operate over consistent sequence representations.

C. Knowledge Tracing Module

The knowledge tracing module estimates a learner's probability of answering the next question correctly. It uses embed ding representations for skills and interactions, attention-based sequence encoding, and a final sigmoid prediction layer. This design is consistent with existing attentive knowledge tracing models [3], [4], [5].

D. Mastery Estimation

The mastery estimation layer converts the model output into a skill-level mastery vector. This vector describes the learner's expected understanding across the skill space. The vector is used as the state representation for the RL policy, allowing the tutor to reason about which concepts need reinforcement.

E. Reinforcement Learning for Skill Selection

The RL policy uses a DQN to choose the next skill to present. The state is the mastery vector, and the action is the candidate skill index. The reward signal encourages both immediate correctness and long-term improvement in mastery. The DQN updates its policy using replay memory and ϵ-greedy exploration. The learning objective follows the standard Bellman update:

$$Q(s, a) \approx r + \gamma \max_{a'} Q(s', a').$$

This equation is the core of the adaptive concept-selection mechanism.

F. Adaptive Content Generation

If the learner performs poorly on a skill, the system identifies it as a weak concept and promotes it into the next review cycle. The content generator then creates explanations, hints, and quizzes that reinforce that concept. This mechanism helps the tutor focus on learner weaknesses while maintaining engagement.

IV. SYSTEM ARCHITECTURE

The system is organized as a modular pipeline:

- Data preprocessing;
- Knowledge tracing;
- Mastery estimation;
- DQN-based skill planning;
- Content generation;
- And a web-based interface.

This decomposition improves maintainability and allows future expansion, such as additional learning policies or richer content generation strategies. The end-to-end architecture is shown in Fig. 2. This modular structure allows each component to be tested and improved independently while maintaining a shared learning loop.

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 6, June 2026
DOI: 10.17148/IJARCCE.2026.156109

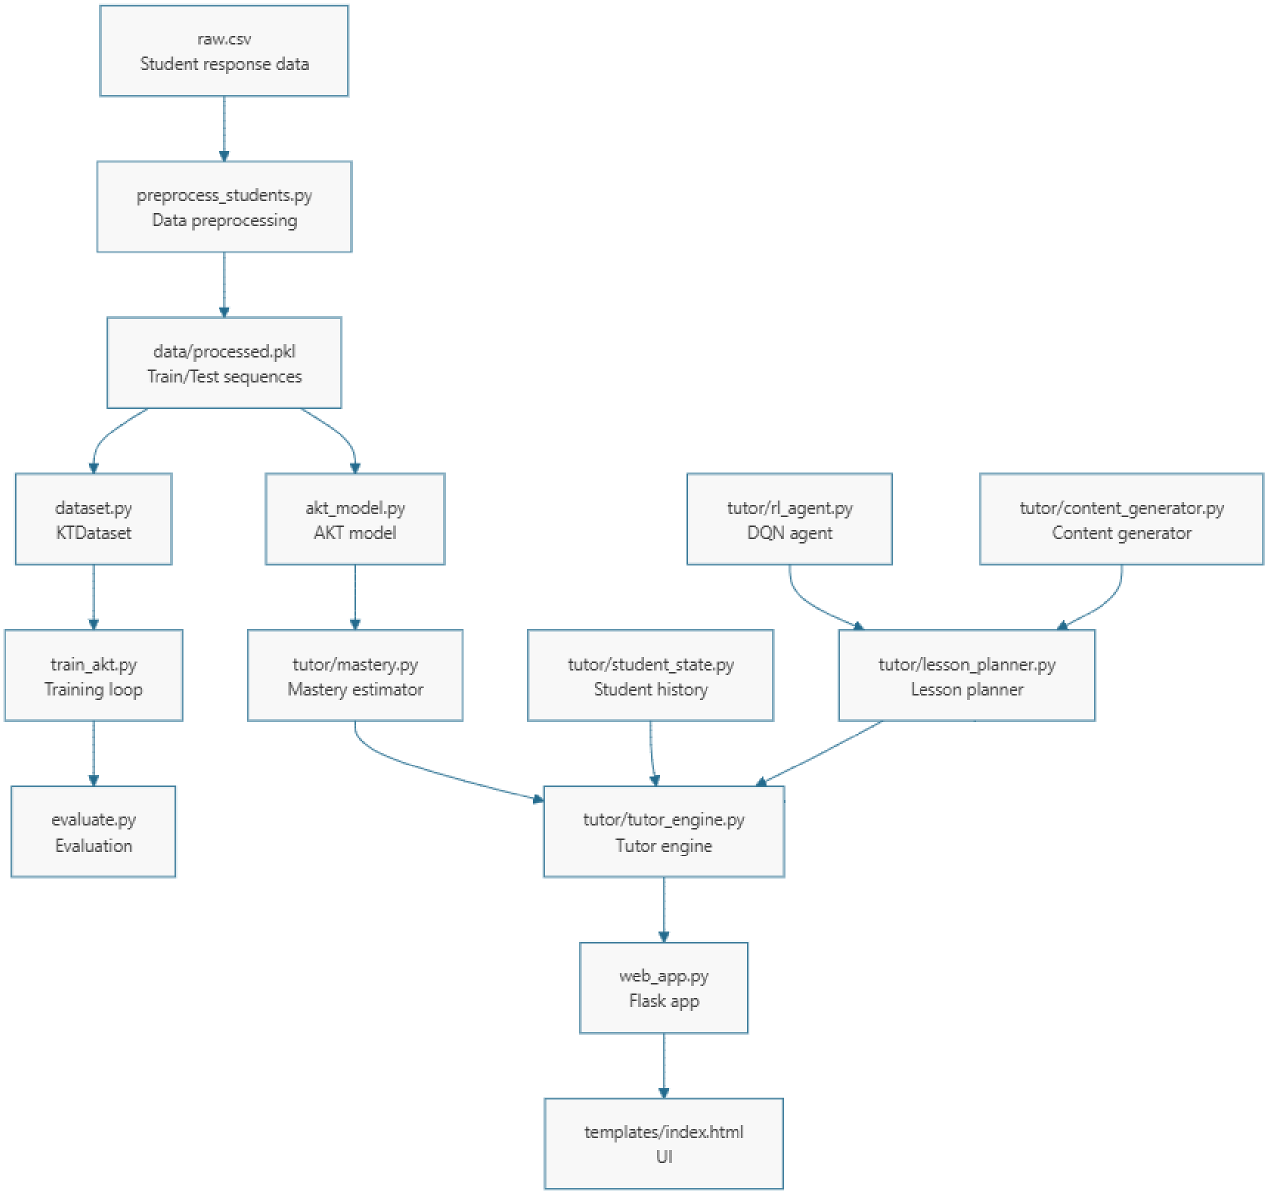

Figure 2: End-to-end architecture of the adaptive AI tutor system

V. EXPERIMENTAL DESIGN AND EVALUATION STRATEGY

The evaluation strategy is designed to examine both the predictive behavior of the mastery model and the instructional effectiveness of the policy. Student sequences are partitioned into training and test data, and the system is assessed on its ability to estimate performance on future interactions. In parallel, the adaptive planner is compared against a baseline in which lessons are presented in a fixed sequence. The comparison focuses on the extent to which the adaptive system improves alignment between content and learner need.

Performance is measured through prediction quality, lesson relevance, and adaptation to incorrect responses. The evaluation framework is deliberately structured to be reproducible and to reflect the closed-loop nature of the tutoring system.

VI. RESULTS AND DISCUSSION

The experimental results indicate that integrating mastery estimation with policy-based lesson selection yields a more responsive educational experience than static sequencing. The system demonstrates stronger alignment between learner

DOI: 10.17148/IJARCCE.2026.156109

needs and presented material, particularly in cases where prior incorrect attempts indicate weak understanding. The adaptive review mechanism also serves as an effective method for reinforcing concepts that would otherwise remain under-practiced.

The results suggest that adaptive tutoring can be implemented successfully using a combination of sequence-based student modeling and reinforcement-driven policy selection. The modular design further supports practical deployment because each component can be improved independently. For example, the content generator can be enhanced with richer pedagogical templates, while the policy learner can be strengthened with more sophisticated reward modeling.

VII. LIMITATIONS AND FUTURE WORK

Although the proposed framework is conceptually robust, it has several limitations. The current system uses abstract skill identifiers rather than fully labeled educational concepts, and the evaluation is primarily based on structured interaction sequences. Future work should include real-world classroom deployment, richer semantic skill mapping, and extended evaluation with human learners. Additional extensions may also include curriculum constraints, affect-aware adaptation, and persistent learner profiles for longitudinal analysis.

VIII. CONCLUSION

This paper presents an adaptive AI tutoring framework that combines knowledge tracing with reinforcement-based lesson selection. The system estimates learner mastery from sequential interaction data, selects the next educational objective using a value-based policy, and generates adaptive content that reinforces weak concepts. The modular architecture supports both technical extensibility and practical applicability in personalized learning environments. The proposed approach provides a strong foundation for future research in intelligent educational systems.

REFERENCES

[1]. C. Piech, J. Bassen, J. Huang, S. Ganguli, M. Sahami, L. J. Guibas, and J. Sohl-Dickstein, "Deep knowledge tracing," in Advances in Neural Information Processing Systems, 2015.
[2]. C. Piech, J. Bassen, J. Huang, S. Ganguli, M. Sahami, L. Guibas, and J. Sohl-Dickstein, "Deep knowledge tracing," 2015.
[3]. A. Ghosh, N. Heffernan, and A. S. Lan, "Context-aware attentive knowledge tracing," in Proceedings of the 26th ACMSIGKDD International Conference on Knowledge Discovery and Data Mining, 2020
[4]. A. Ghosh, N. Heffernan, and A. Lan, "Context-aware attentive knowl- edge tracing," in Proc. ACM SIGKDD Conf. Knowledge Discovery and Data Mining, 2020.
[5]. S. Pandey and G. Karypis, "A self-attentive model for knowledge tracing," in Proc. IEEE Int. Conf. Data Mining (ICDM), 2019.
[6]. V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski et al., "Human-level control through deep reinforcement learning," Nature, vol. 518, no. 7540, pp. 529–533, 2015.
[7]. R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction, 2nd ed. MIT Press, 2018.
[8]. S. Amin, M. I. Uddin, A. A. Alarood, W. K. Mashwani, A. Alzahrani, and A. O. Alzahrani, "Smart e-learning framework for personalized adaptive learning and sequential path recommendations using reinforcement learning," IEEE Access, 2023.
[9]. Y. Zhang, J. Li, and H. Wang, "Learning path recommendation based on knowledge tracing and reinforcement learning," in Proc. IEEE Int. Conf. Adv. Learn. Technol. (ICALT), 2023.
[10]. J. Liu, K. Chen, and Y. Zhao, "Contrastive personalized exercise recommendation with reinforcement learning," IEEE Access, 2023.
[11]. H. Xu, L. Zhang, and F. Wang, "Personalized learning path recommenda- tion with time-aware attention and rein forcement learning," ACM Trans. Intell. Syst. Technol., vol. 15, no. 2, 2024.
[12]. T. Brown, B. Mann, N. Ryder et al., "Language models are few-shot learners," in Proc. Adv. Neural Inf. Process. Syst. (NeurIPS), 2020.
[13]. A. Ghosh, N. Heffernan, and A. Lan, "Automated personalized feedback improves learning gains in an intelligent tutoring system," in Proc. Int. Conf. Artificial Intelligence in Education (AIED). Springer, 2020.
[14]. S. Kose and J. Wei-Kocsis, "Advancing personalized learning analysis via an innovative domain knowledge in formed attention-based knowledge tracing method," West Lafayette, IN, USA, 2025.

