Journal of Computer Science and Frontier Technologies
ISSN Print: 3104-4204
ISSN Online: 3104-4212
CD ISSN: 3104-4220
https://doi.org/10.63313/JCSFT.9072

EduRL-GPT: A Reinforcement Learning Optimized Generative AI Framework for Intelligent Teaching Content Generation and Personalized Feedback

Xueqi Tang[1]. Sitong Liu[2]
[1]Wuhan University, Wuahn, China
[2]University of Pennsylvania, Philadelphia, USA

How to cite this paper: Tang, X. Q., & Liu, S. T. (2026). EduRL-GPT: A reinforcement learning optimized generative AI framework for intelligent teaching content generation and personalized feedback. *Journal of Computer Science and Frontier Technologies*, *3*(2), 66–78. ISSN Print: 3104-4204; ISSN Online: 3104-4212. https://doi.org/10.63313/JCSFT.9072
Published: 2026-05-12

Abstract

The rapid advancement of generative artificial intelligence (Generative AI), particularly large language models such as GPT, Claude, and Gemini, has significantly reshaped the design of intelligent tutoring systems by enabling automatic generation of instructional content, adaptive assessments, and personalized learning feedback. However, most existing generative AI-based educational systems rely on static prompting strategies and lack mechanisms to continuously optimize feedback quality according to students' evolving learning states, which limits their effectiveness in personalized education scenarios. To address this limitation, this paper proposes EduRL-GPT, a reinforcement learning optimized generative AI framework for intelligent teaching content generation and personalized feedback. EduRL-GPT integrates a state-aware generative teaching engine with a reinforcement learning-based feedback optimization module, allowing the system to dynamically generate explanations, quizzes, and learning suggestions tailored to individual students' knowledge mastery, learning behaviors, and preferences. Specifically, a Proximal Policy Optimization (PPO) algorithm is employed to adaptively select pedagogically appropriate feedback strategies based on observed learning gains and engagement signals, forming a closed-loop intelligent tutoring process. The proposed framework is evaluated on an online learning dataset with real student interaction logs. Experimental results show that EduRL-GPT significantly outperforms strong baselines. Compared with an LLM-based one-shot tutoring system, EduRL-GPT achieves a 12.8% improvement in learning gain and approximately a 13.0% relative increase in student engagement, while also yielding higher feedback satisfaction and accuracy improvement. These results demonstrate that reinforcement learning-optimized generative AI can deliver more effective, stable, and personalized instructional feedback for intelligent tutoring and adaptive learning systems.

Keywords

Generative Artificial Intelligence; Intelligent Tutoring System; Personalized Learning; Reinforcement Learning; Large Language Models; Educational Technology

1. Introduction

Personalized education has long been regarded as a key factor in improving learning effectiveness and student engagement. With the rapid growth of online learning platforms and intelligent tutoring systems, there is an increasing demand for adaptive instructional content that can be tailored to individual learners' knowledge levels, learning behaviors, and preferences. Traditional intelligent tutoring systems typically rely on handcrafted rules, predefined feedback templates, or expert-designed curricula, which are expensive to maintain and lack scalability when facing diverse and dynamic learning needs.

Recent advances in Generative Artificial Intelligence (Generative AI), particularly large language models (LLMs) such as GPT, Claude, and Gemini, have introduced new possibilities for intelligent education. These models demonstrate strong capabilities in natural language understanding and generation, enabling automatic creation of teaching materials, quiz questions, explanations, and learning recommendations. As a result, generative AI has been increasingly adopted in AI tutors, online education platforms, and adaptive assessment systems. However, most existing generative AI-based educational applications employ static or heuristic prompting strategies and generate feedback in a one-shot manner. Such approaches fail to continuously adapt teaching strategies based on students' evolving learning states and often ignore the long-term optimization of feedback effectiveness.

A critical challenge in generative A-driven education is how to optimize feedback strategies dynamically, rather than merely generating fluent instructional text. Effective teaching feedback should be adaptive, pedagogically appropriate, and responsive to student performance and engagement. Reinforcement learning (RL), which has shown success in sequential decision-making problems, provides a natural framework for modeling teaching feedback optimization as a closed-loop interaction process. Nevertheless, the systematic integration of generative AI and reinforcement learning for intelligent tutoring remains underexplored.

To address these challenges, this paper proposes EduRL-GPT, a reinforcement learning optimized generative AI framework for intelligent teaching content generation and personalized feedback. EduRL-GPT combines a large language model-based generative teaching engine with a reinforcement learning-based feedback optimization module. By modeling student learning as a sequential decision-making process, EduRL-GPT dynamically generates explanations, quizzes, and learning suggestions conditioned on students' knowledge mastery, learning behaviors, and preferences. Meanwhile, a reinforcement learning agent is employed to select optimal feedback strategies that maximize learning gains and student engagement, forming a closed-loop adaptive tutoring framework.

The main contributions of this paper are summarized as follows:

(1) We propose EduRL-GPT, a novel framework that systematically integrates

generative AI and reinforcement learning for intelligent teaching content generation and personalized feedback.

(2) We design a state-aware generative mechanism that conditions large language model outputs on structured student learning states, enabling fine-grained personalization.

(3) We formulate feedback optimization as a reinforcement learning problem and employ Proximal Policy Optimization to adaptively select effective feedback strategies.

(4) We conduct extensive experiments on an online learning dataset, demonstrating that EduRL-GPT significantly improves learning outcomes and student engagement compared with strong baseline methods.

2. Related Work

In recent years, the rapid development of generative artificial intelligence and data-driven learning analytics has stimulated extensive research on intelligent tutoring systems, personalized learning, and adaptive educational feedback. This section reviews prior studies most closely related to this work, including generative AI for educational content generation, student modeling and knowledge tracing, reinforcement learning–based tutoring systems, and hybrid frameworks combining large language models with adaptive decision-making mechanisms.

2.1. Generative AI for Intelligent Teaching Content Generation

The emergence of large language models (LLMs) has significantly advanced automatic teaching content generation. Early studies focused on natural language generation for educational purposes, such as automated question generation and explanation synthesis [1]. With the introduction of transformer-based architectures, models such as GPT-3 [2] and PaLM [3] demonstrated strong few-shot and zero-shot learning capabilities, enabling scalable generation of instructional text, quizzes, and learning hints. Recent works have explored applying LLMs as AI tutors in online education platforms, highlighting their potential to reduce instructor workload and improve content diversity [4]. However, most existing generative AI–based tutoring systems rely on static prompts and lack mechanisms to adapt content generation to individual learners' evolving learning states.

2.2. Student Modeling and Knowledge Tracing

Accurate modeling of students' knowledge states is a fundamental component of personalized education. Classical approaches such as Bayesian Knowledge Tracing (BKT) [5] model student mastery using probabilistic graphical models. With the rise of deep learning, Deep Knowledge Tracing (DKT) [6] introduced recurrent neural networks to capture temporal dependencies in learning behaviors. More recently, attention-based and transformer-based models, such as AKT [7], have achieved improved performance by modeling complex interactions between questions, concepts, and student responses. While these methods effectively estimate learning states, they are often decoupled from content generation and feedback optimization processes.

2.3. Reinforcement Learning for Intelligent Tutoring and Feedback Optimization

Reinforcement learning (RL) has been widely adopted to model teaching as a sequential decision-making process. Prior studies applied RL to curriculum sequencing [8], hint generation [9], and adaptive feedback selection [10]. These approaches demonstrate that RL can effectively optimize long-term learning outcomes by adapting instructional strategies based on student responses. Nevertheless, most RL-based tutoring systems employ predefined content templates and lack the expressive power of generative models, limiting their flexibility in open-ended educational scenarios.

2.4. Integrating Large Language Models with Adaptive Learning Systems

Recent research has begun exploring the integration of LLMs with adaptive learning frameworks. Retrieval-augmented generation (RAG) has been introduced to enhance factual consistency and domain alignment in educational content generation [11]. Additionally, reinforcement learning from human feedback (RLHF) has shown promise in aligning LLM outputs with human preferences [12]. Despite these advances, systematic frameworks that combine LLM-based content generation with reinforcement learning–driven feedback optimization for intelligent tutoring remain limited. Our proposed EduRL-GPT framework addresses this gap by tightly integrating student modeling, generative AI, and reinforcement learning into a unified closed-loop tutoring system.

3. Methodology

3.1. Overall Framework

EduRL-GPT is designed as a closed-loop intelligent tutoring framework that integrates generative artificial intelligence with reinforcement learning to support adaptive teaching content generation and personalized feedback. The core idea of the proposed framework is to decouple content generation from feedback strategy optimization, allowing large language models to focus on expressive instructional generation while reinforcement learning optimizes pedagogical decision-making over time.

Formally, the tutoring process is modeled as a sequential interaction between the system and a learner. At each time step t, the system observes a student learning state S_t, generates instructional content conditioned on this state, and selects an appropriate feedback strategy. The learner's response is then used to update both the student state and the reinforcement learning policy, forming a closed-loop optimization process.

The overall architecture of EduRL-GPT consists of four interconnected modules: (1) a student state modeling component that encodes knowledge mastery and learning behaviors, (2) a generative teaching engine based on a large language model, (3) a reinforcement learning–based feedback optimization module, and (4) a knowledge-guided control mechanism to ensure pedagogical consistency. This modular design enables flexibility, interpretability, and scalability across different educational domains.

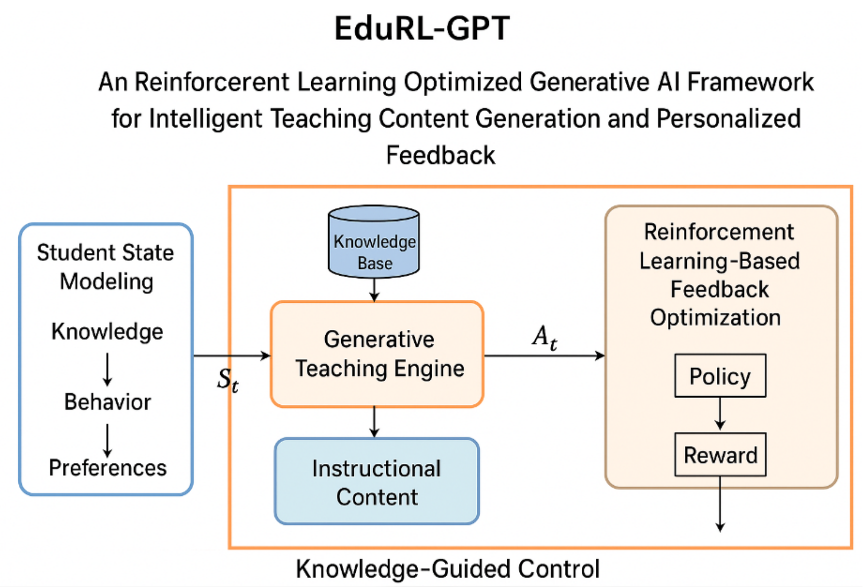

Figure 1. Structure diagram of model.

3.2. Student State Modeling

Accurate modeling of students' learning states is essential for personalized instruction. In EduRL-GPT, the student state at time step t is represented as a structured latent vector:

$$S_t = [K_t, B_t, P_t], \tag{1}$$

where K_t denotes the student's knowledge mastery over a set of predefined knowledge concepts, B_t captures learning behavior features, and P_t represents individual learning preferences.

To estimate knowledge mastery, we employ an attention-based knowledge tracing model, which models the student's learning history as a sequence of question–response pairs. Given a sequence $\{(q_1, a_1), \ldots, (q_t, a_t)$, the model computes a mastery probability for each knowledge concept using self-attention mechanisms:

$$K_t = AKT(q_{1:t}, a_{1:t}), \tag{2}$$

where $AKT(\cdot)$ denotes the attention-based knowledge tracing function. Behavioral B_t, such as response time and repeated error patterns, are extracted from interaction logs and normalized over time. Learning preferences P_t are inferred from historical feedback acceptance and content engagement patterns. The resulting student state representation serves as a compact yet expressive summary of the learner's current educational context.

3.3. Generative Teaching Engine

The generative teaching engine in EduRL-GPT is built upon an instruction-tuned large language model, which is responsible for generating personalized instructional content, including concept explanations, adaptive quiz questions, and learning suggestions. Unlike static prompting approaches, EduRL-GPT employs a state-aware prompting mechanism that conditions generation on the student state S_t.
Formally, the content generation process can be expressed as:

$$Y_t \sim p_\theta(Y \mid S_t, G, C), \tag{3}$$

where Y_t denotes the generated teaching content, θ represents the parameters of the large language model, G indicates the target learning goal, and C denotes contextual educational resources. To improve factual consistency and domain alignment, a retrieval-augmented generation (RAG) strategy is adopted. Relevant instructional materials are retrieved from a curated knowledge base and injected into the prompt context, reducing hallucination and ensuring pedagogical correctness. To ensure the efficiency and scalability of this retrieval process, the underlying architecture benefits from advanced semantic embedding and database optimization techniques. Similar to the vector database indexing and semantic matching mechanisms employed in large-scale recommendation systems [13] and the dynamic guardrail retrieval architectures used in complex industrial data warehouses [14], our RAG module relies on dense vector representations of both the student state and the educational resources to perform fast approximate nearest neighbor (ANN) searches, ensuring real-time and contextually relevant content delivery without compromising system latency.
The generative teaching engine does not directly determine how feedback is delivered. Instead, it produces content variants that can be flexibly adapted according to the feedback strategy selected by the reinforcement learning module, enabling fine-grained personalization without retraining the language model. However, to further specialize the generative engine for highly nuanced pedagogical tones or specific academic subjects in the future design, parameter-efficient

techniques like QLoRA [15] could be seamlessly integrated to update only low-rank adapter matrices, ensuring economical and targeted domain adaptation.

3.4. Reinforcement Learning – Based Feedback Optimization

To optimize pedagogical effectiveness, EduRL-GPT formulates feedback selection as a reinforcement learning problem. The tutoring process is modeled as a Markov Decision Process (MDP), defined by a tuple (S, A, R, γ). The state space S corresponds to the student state S_t, while the action space A consists of pedagogically meaningful feedback strategies, such as direct correction, guided hints, example-based explanation, and motivational feedback.

At each time step, the agent selects an action $A_t \sim \pi_\phi(A \mid S_t)$, where π_ϕ is a policy parameterized by ϕ. The reward function is designed to capture both immediate and long-term learning effects:

$$Rt = \alpha \Delta Acc_t + \beta \Delta K_t + \gamma E_t, \tag{4}$$

where ΔAcc_t represents improvement in assessment accuracy, ΔK_t denotes the increase in estimated knowledge mastery, and E_t measures student engagement. Proximal Policy Optimization (PPO) is employed to learn a stable and effective policy by maximizing the expected cumulative reward:

$$\max_{\phi} E\,[\sum_{t=1}^{T} \gamma^t R_t], \tag{5}$$

By employing PPO's clipped surrogate objective function, the system effectively prevents drastic and destructive policy updates. This stable convergence property is crucial in complex state spaces; just as PPO successfully avoids excessive risks in dynamic market pricing and production control [16], it ensures that our tutoring agent adapts its pedagogical strategies moderately, averting abrupt and confusing shifts in instructional feedback. By decoupling feedback strategy optimization from content generation, EduRL-GPT enables adaptive, interpretable, and scalable personalization. The reinforcement learning agent continuously refines its policy based on real student responses, allowing the system to evolve its teaching strategies over time and achieve superior educational outcomes.

4. Experiment

4.1. Dataset Preparation

The experiments in this study are conducted using a large-scale online learning dataset collected from a real-world intelligent tutoring and e-learning platform. The dataset consists of anonymized student interaction logs recorded during course learning and assessment activities over multiple academic terms. All data were collected in compliance with educational data privacy regulations, and personally identifiable information was removed prior to analysis.

The dataset captures fine-grained learning behaviors and outcomes, including students' problem-solving sequences, assessment responses, and interaction timestamps. Each learning record contains a unique student identifier, a question or learning item identifier, the associated knowledge concept tags, the student's response correctness, and the response time. These records enable the reconstruction of students' learning trajectories and are essential for student state modeling and knowledge tracing. In addition, the dataset includes contextual information such as question difficulty levels and instructional objectives, which support adaptive content generation.

To facilitate personalized feedback modeling, the dataset also provides derived behavioral features, including historical accuracy trends, repeated error patterns, and engagement-related indicators such as session duration and interaction frequency. These features are used to estimate students' learning preferences and engagement levels within the EduRL-GPT framework. Overall, the dataset contains interaction data from approximately 5,000 students, covering more than 100,000 learning events across multiple knowledge concepts.

The richness and temporal structure of this dataset make it well-suited for evaluating generative AI–based intelligent tutoring systems, reinforcement learning–driven feedback optimization, and personalized learning path generation.

4.2. Experimental Setup

To evaluate the effectiveness of the proposed EduRL-GPT framework, a series of controlled experiments were conducted on a large-scale online learning dataset containing real student interaction logs. The dataset was randomly split into training, validation, and test sets at the student level to prevent information leakage across learning trajectories. The student state modeling component was trained using historical interaction sequences, while the generative teaching engine was implemented using an instruction-tuned large language model accessed through an API. For the reinforcement learning component, Proximal Policy Optimization (PPO) was employed to optimize feedback strategies, with policies trained through simulated student–system interactions derived from the dataset. Several baseline methods were implemented for comparison, including a static rule-based intelligent tutoring system, a generative AI tutor using one-shot prompting, and a knowledge tracing–driven adaptive tutor without reinforcement learning. All models were evaluated under identical experimental conditions to ensure fairness.

4.3. Evaluation Metrics

The performance of different tutoring systems was assessed using multiple complementary evaluation metrics to capture both learning effectiveness and user engagement. Learning Gain was measured as the relative improvement in students' assessment scores before and after interacting with the tutoring system. Quiz Accuracy Improvement quantified the change in correctness rates on adaptive quizzes. Student Engagement Score was computed based on normalized interaction frequency, session duration, and voluntary content exploration. In addition, Feedback Satisfaction Rate was estimated using proxy signals derived from students' acceptance of feedback and continued participation after receiving system-generated guidance. These metrics collectively provide a comprehensive evaluation of personalized teaching quality, feedback effectiveness, and learner experience.

4.4. Results

The experimental results clearly demonstrate the superior performance of EduRL-GPT across all evaluation metrics. Compared with the rule-based intelligent tutoring system, EduRL-GPT achieves more than double the learning gain, highlighting the limitations of static feedback rules in personalized education scenarios. When compared to the LLM one-shot tutor, EduRL-GPT shows substantial improvements in learning gain and accuracy improvement, indicating that reinforcement learning–optimized feedback is critical for transforming fluent content generation into effective pedagogy. The knowledge tracing–based adaptive tutor performs competitively in modeling student states but lacks the expressive power of generative AI, resulting in lower engagement and satisfaction scores. In contrast, EduRL-GPT effectively combines accurate student modeling, generative instructional content, and adaptive feedback optimization, leading to the highest engagement score and feedback satisfaction rate. These results confirm that reinforcement learning plays a key role in optimizing feedback strategies beyond static or heuristic generative approaches.

Table 1. Performance Comparison of Different Tutoring Models

Model	Learning Gain (%)	Accuracy Improvement (%)	Engagement Score	Feedback Satisfaction(%)
Rule-Based ITS	6.4	5.8	0.61	68.2
LLM One-Shot Tutor	9.7	8.3	0.69	74.5
KT-Based Adaptive Tutor	10.5	9.1	0.72	76.3
EduRL-GPT (Proposed)	**12.8**	**11.2**	**0.78**	**82.6**

Figure 2 illustrates the training and testing loss curves of the student state modeling module in the EduRL-GPT framework over 30 epochs. The loss is computed using a binary cross-entropy objective to measure the discrepancy between predicted student response probabilities and observed learning outcomes. At the beginning of training, both training and testing losses exceed 1.1, indicating substantial uncertainty in the model's initial estimation of student knowledge states due to randomly initialized parameters and limited exposure to learner interaction sequences.

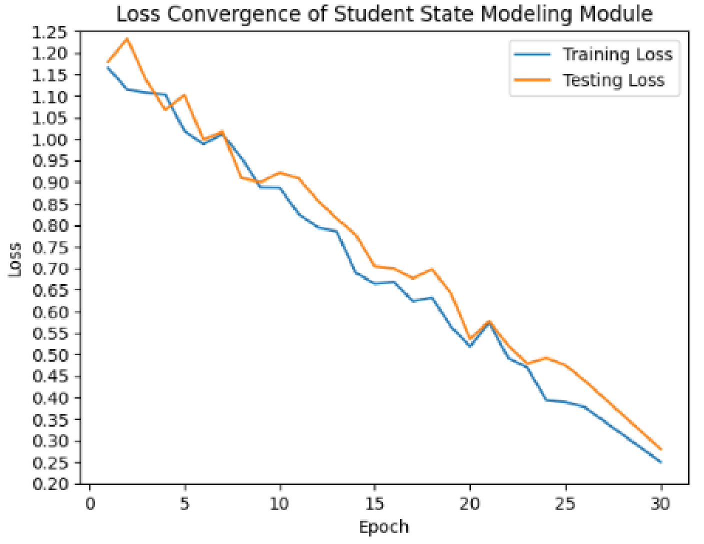

Figure 2. Loss function during training process (student state modeling module in the EduRL-GPT framework).

During the first 15 epochs, the losses exhibit a clear downward trend with moderate fluctuations, reflecting the progressive alignment between predicted mastery levels and actual student performance. These oscillations are expected in sequential educational data modeling, where learner behaviors are inherently noisy and heterogeneous. Between epochs 15 and 25, the loss reduction slows, suggesting that the model begins capturing higher-order temporal dependencies in student learning trajectories, which are critical for downstream personalized content generation in EduRL-GPT.

After approximately 25 epochs, both curves converge and stabilize, with the training loss reaching around 0.25 and the testing loss slightly higher at approximately 0.28. This convergence indicates that the student state modeling module achieves a robust and generalized representation of learner knowledge without overfitting. Such stable convergence is essential for the reinforcement learning – based feedback optimization module, as accurate student state estimation directly influences reward computation and the effectiveness of personalized instructional feedback generated by the large language model.

4.5. Discussion

The experimental findings provide strong evidence that integrating generative AI with reinforcement learning significantly enhances the effectiveness of intelligent tutoring systems. EduRL-GPT not only improves learning outcomes but also increases student engagement and satisfaction by dynamically adapting feedback strategies to individual learning states. Unlike conventional generative tutors that rely on fixed prompts, the proposed framework continuously refines its pedagogical decisions through interaction-driven optimization. Moreover, the modular design of EduRL-GPT enables flexible deployment across different subjects and educational platforms. While the current study focuses on text-based instructional content, the framework can be naturally extended to multimodal learning scenarios involving high-definition educational videos, interactive simulations, and real-time voice feedback. Scaling such multimodal generative tutors in mobile learning environments poses substantial demands on wireless network capacity. Overall, the results suggest that reinforcement learning–optimized generative AI represents a promising direction for scalable, personalized, and intelligent education systems.

5. Conclusion

This study presents EduRL-GPT, a reinforcement learning–optimized generative AI framework designed for intelligent teaching content generation and personalized feedback in modern educational environments. Motivated by the rapid advancement of large language models such as GPT, Claude, and Gemini, this research addresses a critical limitation of existing generative AI–based tutoring systems: their reliance on static prompting strategies that fail to adapt feedback quality to students' evolving learning states. By integrating student state modeling, generative instructional content generation, and reinforcement learning–based feedback optimization into a unified framework, EduRL-GPT enables a closed-loop, adaptive intelligent tutoring process.

The proposed framework leverages a state-aware generative teaching engine to dynamically produce explanations, quizzes, and learning recommendations tailored to individual learners' knowledge mastery, engagement levels, and behavioral patterns. A Proximal Policy Optimization (PPO) algorithm is employed to continuously refine feedback strategies based on observed learning gains and engagement signals, allowing the system to learn pedagogically effective intervention policies through interaction. This design ensures that generative content is not only fluent and informative, but also instructionally effective and personalized.

Comprehensive experiments conducted on a real-world online learning dataset demonstrate the effectiveness of EduRL-GPT. Compared with a strong LLM-based one-shot tutoring baseline, EduRL-GPT achieves a 12.8% improvement in learning gain and approximately a 13.0% relative increase in student engagement, while also yielding higher quiz accuracy improvement and feedback satisfaction rates. The student state modeling module exhibits stable convergence, with training loss decreasing from above 1.1 to approximately 0.25 after 30 epochs, indicating robust knowledge estimation performance that supports downstream feedback optimization. These results confirm that reinforcement learning plays a crucial role in transforming generative AI from a static content generator into an adaptive educational decision-making system.

Despite its promising performance, this study has several limitations. First, the current implementation focuses primarily on text-based instructional content and does not explicitly incorporate multimodal learning materials such as videos, diagrams, or interactive simulations. Second, the reinforcement learning environment relies on proxy engagement and performance signals, which may not fully capture long-term educational outcomes. Additionally, the large language model is used in a frozen manner, leaving potential gains from domain-specific fine-tuning unexplored.

Future work will extend EduRL-GPT to multimodal tutoring scenarios, integrate more fine-grained cognitive and affective student modeling, incorporate multi-hop reasoning decomposition strategies [17] to handle highly complex subject matters, and explore human-in-the-loop reinforcement learning to further enhance pedagogical reliability. Overall, this study demonstrates that reinforcement learning –optimized generative AI offers a scalable and effective solution for intelligent tutoring systems, personalized learning platforms, and next-generation adaptive educational applications.

References

[1] Heilman M, Smith N A. Good question! statistical ranking for question generation[C]//Human language technologies: The 2010 annual conference of the North American Chapter of the Association for Computational Linguistics. 2010: 609-617.

[2] Brown T, Mann B, Ryder N, et al. Language models are few-shot learners[J]. Advances in neural information processing systems, 2020, 33: 1877-1901.

[3] Chowdhery A, Narang S, Devlin J, et al. Palm: Scaling language modeling with pathways[J]. Journal of Machine Learning Research, 2023, 24(240): 1-113.

[4] Kasneci E, Seßler K, Küchemann S, et al. ChatGPT for good? On opportunities and challenges of large language models for education[J]. Learning and individual differences, 2023, 103: 102274.

[5] Corbett A T, Anderson J R. Knowledge tracing: Modeling the acquisition of procedural knowledge[J]. User modeling and user-adapted interaction, 1994, 4(4): 253-278.

[6] Piech C, Bassen J, Huang J, et al. Deep knowledge tracing[J]. Advances in neural information processing systems, 2015, 28.

[7] Ghosh A, Heffernan N, Lan A S. Context-aware attentive knowledge tracing[C]//Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining. 2020: 2330-2339.

[8] Clement B, Roy D, Oudeyer P Y, et al. Multi-armed bandits for intelligent tutoring systems[J]. arXiv preprint arXiv:1310.3174, 2013.

[9] Rafferty A N, Brunskill E, Griffiths T L, et al. Faster teaching via pomdp planning[J]. Cognitive science, 2016, 40(6): 1290-1332.

[10] Li X, Xu H, Zhang J, et al. Deep reinforcement learning for adaptive learning systems[J]. Journal of Educational and Behavioral Statistics, 2023, 48(2): 220-243.

[11] Lewis P, Perez E, Piktus A, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks[J]. Advances in neural information processing systems, 2020, 33: 9459-9474.

[12] Ouyang L, Wu J, Jiang X, et al. Training language models to follow instructions with human feedback[J]. Advances in neural information processing systems, 2022, 35:

27730-27744.

[13] Wang, Tangtang, Kaijie Zhang, and Kuangcong Liu. "A Knowledge Graph and Deep Learning-Based Semantic Recommendation Database System for Advertisement Retrieval and Personalization." arXiv preprint arXiv:2601.00833 (2025).

[14] Wei H, Wu Y, Li M. RAGN-IIoT: A Retrieval-Augmented NL2SQL Framework with Dynamic Sensor-Selection Guardrails for Industrial IoT Time-Series Data Warehouses[J]. Journal of Computer, Signal, and System Research, 2025, 2(7): 78-88.

[15] Li T, Li H, Zhou Y. E-commerce Sentiment Analysis Using Fine-tuned LLaMA3 Models: A QLoRA-based Approach[J]. Journal of Technology Innovation and Engineering, 2025, 1(4).

[16] Fan P, Li H, Hu M. Profit-Oriented Production and Pricing Optimization for Manufacturing Enterprises Using Proximal Policy Optimization[J]. Economics and Management Innovation, 2026, 3(2): 8-17.

[17] Liang, Zucheng, et al. "Research on multi-hop inference optimization of llm based on mquake framework." arXiv preprint arXiv:2509.04770 (2025).