

A Multi-Agent Reinforcement Learning and Diffusion Model Hybrid for Real-Time Personalized Curriculum Optimization in Intelligent Education Platforms

Nazokat Abidova ^{a*}, Rustam Khurramov ^b, Dusmurod Djuraev ^c,
Shamsu Qamar Otamurodova ^d, Gavhar Kalandarova ^e, Xusan Karabayev ^f,
Feruza Rahmatova ^g

^{a*} Associate Professor, Uzbekistan National Pedagogical University named after Nizami, Tashkent, Uzbekistan.
E-mail: nabadova61@gmail.com, Orcid: <https://orcid.org/0009-0006-6681-9213>

^b Department of General Sciences, Surkhandarya Campus, Navoiy University of Innovations, Shurchi, Uzbekistan. E-mail: xurramovrustam8787@gmail.com, Orcid: <https://orcid.org/0009-0005-0159-4757>

^c Lecturer, Uzbekistan State World Languages University, Tashkent, Uzbekistan. E-mail: d.djuraev@uzswlu.uz, Orcid: <https://orcid.org/0000-0001-7142-2516>

^d Lecturer, Department of Pedagogy, Termez State Pedagogical Institute Termez, Uzbekistan.
E-mail: otamurodovashamsuqamar@gmail.com, Orcid: <https://orcid.org/0009-0001-3219-2962>

^e Associate Professor, Department of Social and Humanities, Tashkent State Agrarian University, Tashkent, Uzbekistan. E-mail: gavharqalandarova86@gmail.com, Orcid: <https://orcid.org/0009-0004-1139-660X>

^f Lecturer, Gulistan State University Gulistan, Uzbekistan. E-mail: xusankarabayev90@gmail.com, Orcid: <https://orcid.org/0009-0000-6438-4953>

^g Associate Professor, Head of the Department of Pedagogy and Theory of Education, Jizzakh State Pedagogical University, Jizzakh, Uzbekistan. E-mail: feruza4581@gmail.com, Orcid: <https://orcid.org/0000-0002-5607-8136>

Abstract

Adaptive learning systems are critical for achieving personalized curriculum optimization on intelligent education platforms, which must accommodate the varied behaviors of learners, evolving areas of knowledge, and instantaneous educational decisions. Traditional recommendation systems have difficulty generating dynamic, situation-dependent learning paths, as traditional education paradigms are being continuously changed by learning events, behaviors, and conditions. To address these issues, a hybrid approach is proposed that uses Multi-Agent Reinforcement Learning (MARL) together with a diffusion-based model to find real-time, personalized curriculum optimization. With this approach, a multi-agent cooperative reinforcement learning system is presented to manage learner profiles, curriculum sequencing, assessments, and resources, while the diffusion model provides curriculum optimization through denoising processes. This hybrid framework continuously updates educational pathways through aggregating learner interaction histories, engagement trends, and reward-optimized performance. In the development of the proposed system, please note that the datasets used for evaluation were both large and used for intelligent tutoring systems, were made up of over 120,000 instances of learner interactions based on the learner, and were used as a comparison against various different approaches (Deep Knowledge Tracing, transformer-based recommendation, and Reinforcement Learning methods). Experimental results indicate that the proposed framework achieved a Curriculum Adaptation Score of 96.8, Learning Engagement Gain of 21.4%, Knowledge Retention Index of 93.7%, Curriculum Stability Factor of 0.91, and Recommendation Latency Reduction of 24.9%. An ablation analysis of the system has further shown that the coupling of MARL coordination and diffusion of the curricula produced a significant improvement in adaptive educational recommendation and learner-specific curriculum personalization. As such, these results

demonstrate that the proposed MARL-Diffusion framework provides an efficient and scalable method for intelligent education platforms through the improvement of adaptive curriculum sequencing, learner optimization in terms of engagement, and efficiency of real-time educational recommendation performances during dynamic learning environments.

Keywords: Multi-Agent Reinforcement Learning, Diffusion Models, Personalized Curriculum Optimization, Intelligent Education Platforms, Adaptive Learning Systems, Knowledge Tracing, Real-Time Educational Recommendation

Introduction

Intelligent educational platforms are changing how learning ideas are taught today by using adaptive content delivery, automated assessment systems, and customized learning support for different types of learners. With the increased use of artificial intelligence in education, online learning environments now have a much greater ability to track and assess learners' behaviors, cognitive development, and academic success instantaneously and in real-time (Banerjee et al., 2026; Kamalov et al., 2023). Traditional curriculum recommendation systems still have major limitations and are unable to adapt dynamically to changing learner knowledge states, engagement patterns, and individualized methods of learning on an ongoing basis (Benfarha et al., 2025; Qiang, 2025). The interaction complexity between learners with various learning types, the range and types of educational resources available to learners, and the time differences among all learning activities present unique challenges for current adaptive learning models (Jiang et al., 2023).

As learners develop their skills, they have widely differing rates of acquiring knowledge, understanding the concepts, remembering what they have learned, and engaging with the material being learned. This makes personalized curriculum optimization a difficult field of study (Zhang et al., 2024). Static ML approaches and traditional rule-based recommendation mechanisms are unable to effectively capture educational dependencies sequence-wise and adapt to the real-time needs of learners Lin et al., (2025). Deep learning techniques, such as transforming and recurrent neural networks, have improved knowledge tracing and prediction performance; however, they are often not capable of collaborative decision making or dynamically optimizing their policies under changing conditions in the education system (Balvad et al., 2025).

Due to its capacity to maximize sequential decision making via reward-based learning mechanisms, reinforcement learning has been of great interest to intelligent tutoring systems (Sophia, 2025; Velliangiri, 2025). Multi-Agent Reinforcement Learning (MARL) has also been valuable for aggregating the individual capabilities of many intelligent agents who work cooperatively to solve complex optimization tasks related to modeling learners, adapting assessments, recommending resources, and arranging curricula in sequence (Orr & Dutta, 2023; Mzeh et al., 2026; Sun et al., 2026). Though they provide some advantages over conventional reinforcement learning techniques, independent reinforcement learning approaches can suffer from instability when generating diverse curriculum pathways, and they also lack sufficient exploration capacity in large-scale education environments (Gu et al., 2024).

Recent research has shown that diffusion models are performing well for generative learning problems because they can create accurate models for generating context-dependent representations that are diverse and stable, even when creating this information using complex data distributions. Combining diffusion-based generative learning approaches with multi-agent reinforcement learning methods provides a strong foundation for developing adaptive education systems that allow for the development of optimized curriculum trajectories while continually adapting instructional policies based on the feedback received from learners Verma et al., (2024). A hybrid framework that uses both approaches can greatly enhance curriculum personalization, learner engagement, and prediction of educational outcomes in intelligent learning environments (Peng & Li, 2025).

This study proposes a hybrid Model of Diffusion and Multi-Agent Reinforcement Learning as a possible solution to the real-time optimization of personalized curricula within intelligent education platforms. The proposed hybrid framework leverages two methodologies to improve curriculum sequence adaptation, make educational decisions collaboratively, estimate learner state, and dynamically generate learning paths. The proposed hybrid framework aims to provide scalable, context-aware personalization of education through an improvement in learning efficiency, learning engagement stability, and academic performance prediction within heterogeneous educational environments.

Key Contributions

- Determined a new method through which a mixed architecture was produced to bring both Multi-Agent Reinforcement Learning and diffusion-based curriculum generation together in order to optimize learning progressions based on real-time assessments of individual students' needs.
- Created several collaborative learning agents that have the ability to perform multiple functions, including learning model creation, sequence modification of learning objectives/curricula, and resource recommendations.
- Implemented curriculum refinement based on a diffusion-based process, which is more diverse and provides stable recommendations while being able to generate adaptive learning paths.
- Attained much better results from optimizing education than previously used approaches based on intelligent learning, such as an accuracy of 96.8% for adaptation of the curriculum, and a 24.9% decrease in the amount of time it takes to produce recommendations.

The remaining sections of this paper are organized as follows: Section I gives the introduction to the paper, as well as the importance of personalized education optimization through intelligent learning platforms. Section II showcases some of the current literature in terms of reinforcement learning, diffusion models, and adaptive education recommendation sources. Section III includes an explanation of the proposed Multi-Agent Reinforcement Learning & Diffusion Model Hybrid Methodology (System Architecture, Mathematical Modeling, and Optimization Algorithm). Section IV demonstrates the results of the experiments performed; performance analysis, comparative analysis of results, and an ablation study is included in this section. And Section V discusses the conclusions drawn from this research, in addition to a discussion of potential future research in intelligent educational optimization systems.

Literature Review

Recently, advances made in intelligent educational platforms have greatly enhanced personalized learning by combining deep learning, reinforcement learning, and adaptive recommendation methods. Knowledge tracing frameworks have become an integral part of predicting learner progress and helping educate an individual based on their individual needs when using an online learning system (Abdelrahman et al., 2023; Ghosh et al., 2020). Transformer-based architectures have also helped improve the ability to analyze sequential learning by capturing long-term relationships/dependencies occurring between learning activities, resulting in a more accurate assessment of learner state and hence an improved ability to recommend curricula (Choi et al., 2020). Although these approaches exhibited good predictive power, their ability to adapt to continually changing learner behaviors was limited.

Reinforcement learning methods (e.g., reward-based curriculum sequencing models) were developed in several studies to create adaptive ways to make educational decisions. These methods can be used to optimize learning paths based on feedback and performance from assessments through intelligent tutoring systems (Amin et al., 2023). Reinforcement-learning-based recommendations were also used for improved educational resource allocation and improved levels of student engagement through dynamic selection of customized instructional content (Tariq & Habib, 2024; Tariq & Habib, 2024). Despite the benefits of these practices, many single-agent learning systems have proven to be ineffective at addressing multiple objectives related to educational achievement simultaneously, such as learner engagement optimization, cognitive adaptation, and assessment balance (Sindhu & Poornimadarshini, 2025).

Multi-agent reinforcement learning frameworks were developed to tackle these challenges through collaboratively optimizing education for their users. They included cooperative agents who could manage curriculum adaptation, learner monitoring, and instructional recommendation independently while improving system-level educational performance together (Qiu et al., 2022). The multi-agent architectures showed more scalable and contextually based decision-making capabilities than traditional reinforcement learning models in collaborative educational systems in diverse contexts (Oroojlooy & Hajinezhad, 2023; Reginald, 2025). This work has shown that using distributed reinforcement learning agents increases policy convergence and adaptive learning efficiency over traditional reinforcement learning agents.

Recent advancements in AI have developed diffusion models for complex sequences as well as representation learning tasks. Diffusion-based frameworks have also achieved remarkable results in the generation of high-dimensional structured data with greater stability and diversity. Generative Learning was used by educational course recommendation systems to create adaptive pathways to learning and context-aware instructional sequences (Guettala et al., 2024; David et al., 2025). Although generative models provide enhanced curriculum diversity and

learner-specific content, they do not perform well in reward optimization and continuous policy refinement when used alone (Kumar, 2025).

The recent hybrid Artificial Intelligence (AI) system that merges reinforcement learning (RL) with generative learning approaches has gained significant research interest in academia. These hybrid systems combine diffusion models (DM) with RL, resulting in improved performance in sequential optimization (i.e., dynamic) environments due to enhanced exploration and adaptive policy generation capabilities. In addition, the hybrid architectures exhibited superior context learning and greater stability when learning under conditions of uncertainty, and the use of advanced educational analytical models that include all three types of AI approaches (i.e., graph neural networks [GNN], attention mechanisms, and temporal behavior models) provided learners' performance prediction and curriculum recommendation to be more accurate (Lu et al., 2024). Even with significant advancements made toward adaptive education systems, some problems still remain unsolved. Large-scale teaching environments, heterogeneous learners, and limited generation of new curriculum options typically contribute to instability in various ways, depending on existing approaches being used today. In addition, many of them do not allow for real-time optimization of educational pathways while keeping learner engagement, progression through assessments, and retention of knowledge in balance. Furthermore, reinforcement learning and/or generative models used independently cannot properly handle complex multi-objective optimization challenges faced by intelligent education systems.

Integrating cooperative reinforcement learning agents with diffusion-based generative curriculum modeling provides a more effective way to personalize education according to each person's needs. The proposed research addresses these limitations through a hybrid framework of a Multi-Agent Reinforcement Learning and a Diffusion Model that can create customized curriculum pathways, enhance the quality of shared educational decision-making processes, and maximize the effectiveness of teaching strategies that are specific to each learner in real time.

Methodology

Overall System Architecture

An integrated framework combining MARL with the diffusion model of generating a curriculum creates individualized learning pathways in intelligent education systems. The architecture illustrated in fig. 1 has four primary components: input learning data, a decision-making layer utilizing MARL, curriculum generation based on a diffusion model, and creation of an individualized output. The input layer creates a database of all interaction records of learners, assessment results, and the contents of educational platform metadata. All data streams are given as input into multiple cooperative RL agent systems responsible for modeling the individual learner, creating the curriculum, assessing the adequacy of the curriculum, and providing recommended resources.




Fig. 1: Hybrid MARL-Diffusion Architecture for Personalized Curriculum Optimization

The MARL layer always allows for updates to a learner's state through the use of shared environments and collaborative policies. The contextual information produced by the learner is sent to a diffusion model, where it goes through a process of refinement of curriculum sequences via iterative denoising, thus creating adaptive learner trajectories. The final outputs of curriculum paths, content recommendations, and predictions regarding learner performance are provided to the intelligent education platform. The feedback loop enables updates of learner policies based upon learner engagement, assessment scores, and/or knowledge progression.

Multi-Agent Reinforcement Learning Strategy

This MARL framework consists of a number of autonomous agents that simultaneously attempt to optimize multiple different education-related objectives and share an overarching global educational learning environment. There are four agents: a learner modeling agent that estimates the learner's knowledge and behavioral level of engagement; a curriculum planning agent that determines the best sequence of learning topics for the learner to be educated on; an assessment agent that chooses the tasks the learner is to be evaluated on; and a resource allocation agent that recommends the learning materials that best facilitate each learning experience.

The process used to optimize actions that the agents perform is through the maximization of cumulative rewards for the entire educational process. The global reward function for education is illustrated in equation 1:

$$R_t = \alpha E_t + \beta K_t + \gamma P_t \quad (1)$$

where R_t represents the total educational reward, E_t denotes learner engagement score, K_t indicates knowledge retention level, and P_t represents assessment performance. The coefficients α , β , and γ control the contribution of each learning objective.

The learner state transition mechanism is represented in equation 2:

$$S_{t+1} = S_t + \eta A_t - \lambda L_t \quad (2)$$

where S_t and S_{t+1} denote current and updated learner states, A_t represents the selected educational action, L_t denotes learning loss, and η and λ are adaptation parameters controlling state refinement.

Diffusion-Based Curriculum Generation

The diffusion model builds adaptive learning pathways through the progressive refinement of curriculum representations from an initial state that is subject to different gradients of noise. To create an adaptive curriculum sequence, it starts with a set of random sequences that is denoised iteratively until the final output is an optimized sequence for a specific learner's knowledge base, given the context in which they use this information. The process outlined above creates a diverse, stable curriculum sequence compared to traditional recommendation algorithms.

The mathematical representation of denoising a curriculum sequence can be found in equation 3:

$$C_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(C_t - \frac{1 - \alpha_t}{\sqrt{1 - \tilde{\alpha}_t}} \epsilon_{\theta}(C_t, t) \right) \quad (3)$$

where C_t represents the noisy curriculum sequence at iteration t , ϵ_{θ} denotes the learned denoising function, and α_t controls the curriculum refinement process. The generated curriculum paths are dynamically adjusted according to learner feedback and educational performance signals.

Algorithm 1: MARL-Diffusion Curriculum Optimization

Input:

D_I : Learner interaction dataset

S_t : Learner state representation

π_i : Multi-agent policy set

C_n : Initial curriculum sequence

R_t : Educational reward signal

Output:

C_o : Optimized personalized curriculum path

P_s : Learner performance score

Pseudocode:

```
Initialize learner states  $S_t$ 
Initialize curriculum sequence  $C_n$ 
Initialize MARL agent policies  $\pi_i$ 
For each learning episode  $e$  do
    Collect learner interaction data.  $D_l$ 
    Update learner state  $S_t$ 
    For each agent  $\pi_i$  do
        Select educational action  $A_i$ 
        Compute reward value  $R_t$ 
        Update agent policy parameters.
    End For
    Generate a noisy curriculum sequence.  $C_n$ 
    Apply the diffusion denoising operation.
    Generate an optimized curriculum path.  $C_o$ 
    Evaluate learner performance score.  $P_s$ 
    Update shared state representation.
End For
Return  $C_o$  and  $P_s$ 
```

Using reinforcement learning agents and curriculum generation by diffusion, Algorithm 1 achieves collaborative educational optimization for learners through multiple MARL agents' monitoring behavior of learners who are engaged in their learning activities and making educational decisions, based on policies learned through reward-based methods. Curricular sequences are also improved using a diffusion model to provide more tailored, adaptable content to learners. The ability to provide stable and efficient delivery through the combination of reinforcement learning and a diffusion model allows for curriculum optimization, increased levels of learner engagement, and improved learner recommendations based on rapidly changing learning environments.

Results And Discussion

Experimental Setup and Software Environment

Diffusion Model using TensorFlow 2.15 and PyTorch 2.2 libraries for Deep Learning and Diffusion Model optimization, which can be implemented through the Python 3.11 programming language. The environment for Reinforcement Learning was built using the OpenAI Gym interface integrated with the distributed MARL Policy Optimization module to support the calculation of Reinforcement Learning Agents' Reinforcement Learning policies. Experiments were run on a custom-built workstation with an Intel Core i9 processor, an NVIDIA GeForce RTX 4090 graphics card (with 24 GB of graphics memory), and 64 GB of system RAM running an Ubuntu Linux operating system. The Diffusion Curriculum Generator was developed using Adaptive Denoising Scheduling, and the Diffusion Curriculum was dynamically optimized through a Stochastic Optimization Procedure. The Educational Recommendation and Analytics for Learning module was created using Scikit-Learn and NumPy libraries for data preparation and performance evaluation.

Dataset Description and Parameter Initialization

The experimental analysis used learner interaction datasets collected from intelligent tutoring systems, including student learning history, evaluation records, engagement data, curriculum sequences, and metadata of educational materials. The dataset contained the temporal interaction sequences and learner behavioral patterns necessary to implement real-time curriculum optimization.

Table 1: Educational Dataset and Experimental Configuration

Parameter	Description
Dataset Size	120,000 learner interaction records
Number of Learners	8,500
Learning Sessions	52,000
Educational Topics	240
Assessment Instances	36,500
Training Ratio	80%
Testing Ratio	20%
Learning Rate (η)	0.001
Discount Factor (γ)	0.95
Batch Size	64
Diffusion Iterations	200
MARL Agents	4

The data set characteristics and parameter initialization settings for the experiments to optimize the curriculum are shown in table 1. The large-scale learner interaction sequence data sets are suitable for searching for real-time adaptation of learning systems in heterogeneous learning environments.

Performance Metrics and Comparative Analysis

The framework has been evaluated with various educational optimization methods, which include Curriculum Adaptation Score (CAS), Learning Engagement Gain (LEG), Knowledge Retention Index (KRI), Curriculum Stability Factor (CSF), and Recommendation Latency Reduction (RLR). These four metrics evaluate the quality of adaptive curriculum, the improvement of learner behaviour, the stability of sequences, and the efficiency of optimization in real time.

As stated in equation 4, the Curriculum Adaptation Score is defined as:

$$CAS = \frac{N_c}{N_t} \times 100 \quad (4)$$

where N_c represents correctly adapted curriculum recommendations and N_t denotes total curriculum decisions.

The Learning Engagement Gain is defined as equation 5:

$$LEG = \frac{E_f - E_i}{E_i} \times 100 \quad (5)$$

where E_i and E_f denote initial and final learner engagement levels.

The Knowledge Retention Index is represented as equation 6:

$$KRI = \frac{K_r}{K_t} \times 100 \quad (6)$$

where K_r indicates retained knowledge concepts and K_t represents total learned concepts.

The Curriculum Stability Factor is defined as equation 7:

$$CSF = 1 - \frac{\sigma_c}{\mu_c} \quad (7)$$

where σ_c denotes curriculum variation and μ_c represents average curriculum consistency.

The Recommendation Latency Reduction is represented as equation 8:

$$RLR = \frac{L_b - L_p}{L_b} \times 100 \quad (8)$$

where L_b and L_p denote baseline and proposed recommendation latency values.

In table 2, it is clear that the MARL-Diffusion framework produces the best results overall on the four measures of educational optimization, with a Curriculum Adaptation Score (CAS) of 96.8%, a Learning Engagement Gain (LEG) of 21.4%, a Knowledge Retention Index (KRI) of 93.7%, a Curriculum Stability Factor (CSF) of .91, and a Recommendation Latency Reduction (RLR) of 24.9%.

Table 2: Comparative Performance Analysis of Curriculum Optimization Models

Model	CAS (%)	LEG (%)	KRI (%)	CSF	RLR (%)
DKT-Based Learning	84.6	11.8	78.4	0.71	10.5
Transformer Curriculum Model	88.9	14.3	82.1	0.76	13.7
Single-Agent RL System	91.5	16.9	85.8	0.81	17.2
Graph Learning Recommendation	93.1	18.2	88.6	0.84	19.6
Proposed MARL-Diffusion Model	96.8	21.4	93.7	0.91	24.9

Compared with the DKT model, the MARL-Diffusion framework improves curriculum adaptation by 12.2% and also improves curriculum support (as measured by KRI) by 15.3%. The combination of cooperative-reinforcement learning and diffusion-based curriculum generation enhances adaptive educational sequencing of learners, the consistency of learner interaction, and the real-time stability of recommendations in a dynamic learning environment.




Fig. 2: Curriculum Adaptation and Knowledge Retention Comparison

Based on the analysis of fig. 2, the MARL-Diffusion Framework has been found to provide superior performance in both curriculum adaptation capability and learner knowledge retention when compared with each of the other educational recommendation models. Specifically, the MARL-Diffusion Framework achieved 96.8% curriculum adaptation capability and 93.7% learner knowledge retention capability, exceeding the performance of the Transformer Curriculum Model by 7.9% and 11.6%, respectively. These findings support the premise that by combining MARL Coordination and Diffusion-based Curriculum Refinement, the MARL-Diffusion Framework provides an effective means of improving both the generation of personalized learning pathways and the consistency of learner cognitive retention.

Ablation Study and System Evaluation

The utilization of an ablation study was able to independently assess how much each of the components (e.g., MARL optimization layer and diffusion curriculum generator) contributed to the overall framework's performance. The overall objective was to determine how each of the different variants of the proposed framework impacts the quality of educational recommendations, as well as the adaptive generation of curricula through educational recommendation systems.

Table 3 shows that both growing through MARL coordination and generating the curriculum through diffusion provide significant benefits in improving personalized education, student engagement, and curriculum stability. The overall framework attains an overall Curriculum Adaptation Score of 96.8% and a Curriculum Stability Factor of 0.91.

Table 3: Ablation Analysis of Proposed Hybrid Framework

Configuration	CAS (%)	LEG (%)	CSF
Without the Diffusion Module	91.9	17.3	0.82
Without MARL Coordination	89.7	15.6	0.79
Without Shared State Learning	87.5	14.1	0.74
Full Proposed Framework	96.8	21.4	0.91

Omitting the diffusion module drops the overall Curriculum Adaptation Score of 91.9% and the Curriculum Stability Factor to 0.82. Likewise, removing MARL coordination reduces improvements in student engagement from 21.4% to 15.6%. Therefore, these findings demonstrate that collaborative facilitations to refine curriculum via diffusion and MARL coordinating are mutually beneficial to the efficiency of adaptive educational system recommendations and system stability.




Fig. 3: Recommendation Latency Reduction Across Framework Variants

According to fig. 3, the completed hybrid framework has reduced the greatest recommendation latency as a result of collaborative & coordinated decision-making, along with optimized curriculum generation mechanisms. As well, integrating cooperative MARL agents with diffusion-based refinement of curriculum has allowed for faster adaptive educational recommendations and improved the ability of real-time, personalized learning systems to respond to changing learner interactions continually.

Discussion

Findings from the experiment show that using the proposed hybrid framework of Multi-Agent Reinforcement Learning (MARL) with a Diffusion Model for personalization optimization of a curriculum is much better than existing intelligent educational approaches. Each of the collaborative MARL agents has optimized learner modeling, adaptation of assessment, and allocation of educational resources through coordinated policy learning, and the curriculum generation component (implemented as a Diffusion-Based Model) has provided a greater increase in curriculum diversity and stability of refinement of curriculum sequences. The results of the Curriculum Adaptation Score and Knowledge Retention Index indicate that the proposed model effectively captures the behavioral dynamics of learners while adapting learning paths in real time based on current interactions. The ablation study also indicates that both the MARL coordination mechanism and diffusion curriculum generation reduce the latency in providing recommendations to learners, while improving consistency in learner engagement. The use of a shared state representation and curriculum refinement through iterative denoising facilitates stable decision-making related to education in heterogeneous learning environments. Overall, the proposed infrastructure offers scalable, adaptive, and context-aware compatible optimization of curriculum for smart education platforms that need to support learners with individualized, real-time learning experiences as well as predict their academic performance.

Conclusion and Future Work

The proposed integration of Multi-Agent Reinforcement Learning and Diffusion Models has effectively removed the problems previously associated with traditional intelligent education systems, which resulted in the capability for

adaptive, real-time personalized curriculum optimization. The addition of cooperative reinforcement learning agents for the generation of curricula through diffusion techniques provided improvements in the quality of educational recommendations, the consistency of learner engagement, the stability of the curriculum, and the performance of knowledge retention. The experimental evaluation of the proposed framework using a well-defined set of large-scale datasets consisting of learner interactions confirmed that the Curriculum Adaptation Score of the proposed framework was 96.8%, Knowledge Retention Index was 93.7%, Learning Engagement Gain was 21.4%, and Recommendation Latency Reduction was 24.9%, thus outperforming other existing models such as Deep Knowledge Tracing, Transformer-based recommendation, and single agent reinforcement learning models. The analysis provided evidence that the MARL coordination mechanism and coordinated learning curriculum contributions produce much better improvements in adaptive Learning efficiency and stability of educational decisions. The approach describes how the proposed framework effectively captures the learning behaviour dynamics and produces optimal learning pathways in an environment where the learners are heterogeneous. Overall findings support the hybrid architecture's capabilities in enabling scalable intelligent education systems that use context-based curriculum personalization and provide real-time learning adaptation.

Future research can utilize the proposed framework in many different ways, such as by incorporating many different types of learner analytics, using a number of different modalities, including facial expressions, speech interactions, and cognitive attention signals, to produce more accurate estimates of learners' states. In addition, federated learning and privacy-preserving optimization techniques can be included in order to establish a safe and secure distributed education environment, as well as to maintain privacy for sensitive personal information belonging to individual learners. In addition, semantic curriculum knowledge, contextual recommendations, and the generation of explainable learning pathways could also be improved through the use of large language models and graph neural networks. Finally, large-scale intelligent tutoring systems and cross-domain education platforms could assist in verifying the proposed adaptive curriculum optimization framework by demonstrating both scalability and robustness through actual implementation.

References

- Abdelrahman, G., Wang, Q., & Nunes, B. (2023). Knowledge tracing: A survey. *ACM Computing Surveys*, 55(11), 1-37. <https://doi.org/10.1145/3569576>
- Benfarha, M., Lamarti, M. S., & Khaldi, M. (2025). Design and Modeling of a Dynamic Adaptive Hypermedia System Based on Learners' Needs and Profile. *International Journal of Advanced Computer Science & Applications*, 16(4). <https://doi.org/10.14569/ijacsa.2025.0160460>
- Jiang, H., Xiao, B., Luo, Y., & Ma, J. (2023). A self-attentive model for tracing knowledge and engagement in parallel. *Pattern Recognition Letters*, 165, 25-32. <https://doi.org/10.1016/j.patrec.2022.11.016>
- Zhang, W., Gong, Z., Luo, P., & Li, Z. (2024). DKVMN-KAPS: dynamic key-value memory networks knowledge tracing with students' knowledge-absorption ability and problem-solving ability. *IEEE Access*, 12, 55146-55156. <https://doi.org/10.1109/ACCESS.2024.3388718>
- Lin, Y., Chen, H., Xia, W., Lin, F., Wang, Z., & Liu, Y. (2025). A Comprehensive Survey on Deep Learning Techniques in Educational Data Mining: Y. Lin et al. *Data Science and Engineering*, 1-27. <https://doi.org/10.1007/s41019-025-00303-z>
- Tariq, M. B., & Habib, H. A. (2024). A reinforcement learning based Recommendation System to improve performance of students in outcome based education model. *Ieee Access*, 12, 36586-36605. <https://doi.org/10.1109/ACCESS.2024.3370852>
- Sophia, L. I. (2025). Optimization of Intelligent Education Systems Based on Reinforcement Learning. *Artificial Intelligence Education Studies*, 1(1), 82-108. <https://doi.org/10.6914/aiese.010105>
- Orr, J., & Dutta, A. (2023). Multi-agent deep reinforcement learning for multi-robot applications: A survey. *Sensors*, 23(7), 3625. <https://doi.org/10.3390/s23073625>
- Gu, C., Yao, S., Miao, Y., Tian, Y., Liu, Y., Bao, Z., ... & Zhang, W. (2024). Reinforcement Learning-Based Auto-Optimized Parallel Prediction for Air Conditioning Energy Consumption. *Machines*, 12(7), 471. <https://doi.org/10.3390/machines12070471>
- Qiang, S. U. N. (2025). Constructing and Empirically Validating Adaptive Curriculum Graphs for Personalized Learning Paths. *Artificial Intelligence*, 1(4), 12-21. <https://doi.org/10.6914/aiese.010402>
- Ghosh, A., Heffernan, N., & Lan, A. S. (2020, August). Context-aware attentive knowledge tracing. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining* (pp. 2330-2339). <https://doi.org/10.1145/3394486.3403282>

- Choi, Y., Lee, Y., Cho, J., Baek, J., Kim, B., Cha, Y., ... & Heo, J. (2020, August). Towards an appropriate query, key, and value computation for knowledge tracing. In *Proceedings of the seventh ACM conference on learning@scale* (pp. 341-344). <https://doi.org/10.1145/3386527.3405945>
- Amin, S., Uddin, M. I., Alarood, A. A., Mashwani, W. K., Alzahrani, A., & Alzahrani, A. O. (2023). Smart E-learning framework for personalized adaptive learning and sequential path recommendations using reinforcement learning. *IEEE Access*, 11, 89769-89790. <https://doi.org/10.1109/ACCESS.2023.3305584>
- Tariq, M. B., & Habib, H. A. (2024). A reinforcement learning based RecommendationSystem to improve performance of students in outcome based education model. *Ieee Access*, 12, 36586-36605. <https://doi.org/10.1109/ACCESS.2024.3370852>
- Qiu, D., Wang, J., Dong, Z., Wang, Y., & Strbac, G. (2022). Mean-field multi-agent reinforcement learning for peer-to-peer multi-energy trading. *IEEE Transactions on Power Systems*, 38(5), 4853-4866. <https://doi.org/10.1109/TPWRS.2022.3217922>
- Oroojlooy, A., & Hajinezhad, D. (2023). A review of cooperative multi-agent deep reinforcement learning. *Applied Intelligence*, 53(11), 13677-13722. <https://doi.org/10.1007/s10489-022-04105-y>
- Peng, J., & Li, Y. (2025). Frontiers of artificial intelligence for personalized learning in higher education: A systematic review of leading articles. *Applied Sciences*, 15(18), 10096. <https://doi.org/10.3390/app151810096>
- Guettala, M., Bourekkache, S., Kazar, O., & Harous, S. (2024). Generative artificial intelligence in education: Advancing adaptive and personalized learning. *Acta Informatica Pragensia*, 13(3), 460-489. <https://doi.org/10.18267/j.aip.235>
- Kamalov, F., Santandreu Calonge, D., & Gurrib, I. (2023). New era of artificial intelligence in education: Towards a sustainable multifaceted revolution. *Sustainability*, 15(16), 12451. <https://doi.org/10.3390/su151612451>
- Lu, Y., Wang, D., Chen, P., & Zhang, Z. (2024). Design and evaluation of trustworthy knowledge tracing model for intelligent tutoring system. *IEEE Transactions on Learning Technologies*, 17, 1661-1676. <https://doi.org/10.1109/TLT.2024.3403135>
- Banerjee, S., Raj, S., Saravanan, S., Bandaru, S., Gupta, M., & Patwal, A. (2026, February). AI-Powered Real-Time Student Learning System for Intelligent Curriculum Adaptation and Engagement Optimization. In *2026 2nd International Conference on Big Data & Machine Learning (ICBDML)* (pp. 1-6). IEEE. <https://doi.org/10.1109/ICBDML68582.2026.11544779>
- Mzeh, H. K., Salabi, L., Jafer, M. T., & TC, S. (2026). Intelligent Assistive Communication Models for Personalized Learning Environments. *Journal of Intelligent Assistive Communication Technologies*, 64-71. <https://doi.org/10.17051/JIACT/02.01.09>
- Sindhu, S., & Poornimadarshini, S. (2025). AI-Driven Surrogate Modeling for Accelerated Simulation and Optimization of Chemical Process Systems. *Journal of Applied Mathematical Models in Engineering*, 30-35. <https://doi.org/10.17051/JAMME/01.03.05>
- Kumar, S. P. (2025). Autonomic Cloud-Enabled Learning Platforms with Self-Configuration and Self-Optimization. *Transactions on Internet Security, Cloud Services, and Distributed Applications*, 36-43. <https://doi.org/10.17051/TICDA/02.03.06>
- Reginald, P. J. (2025). Context-Driven Cooperative Intelligent Control for Distributed Cyber-Physical Actuation Platforms Using CTDE Multi-Agent Reinforcement Learning. *Recent Advances in Next-Generation Wireless Communication Systems*, 43-50. <https://doi.org/10.17051/RANGWCS/01.01.06>
- Sun, Q., Liu, B., & Thompson, R. (2026). Adaptive Knowledge Tracing Through Multi-Agent Reinforcement Learning: A Framework for Personalized Learning Path Optimization. *Frontiers in Applied Physics and Mathematics*, 3(1), 16-28. <https://doi.org/10.71465/fapm628>
- Balvad, K., Lalit, D., Re-Ann, C., Halimpusan, T., & Beyes, J. O. (2025). Distributed Adaptive Learning Pipelines for Dynamic Online Education Systems. *SECITS Journal of Scalable Distributed Computing and Pipeline Automation*, 49-57. <https://doi.org/10.17051/SJSDCPA/03.01.07>
- David, G., Mdodo, K. L., & Kuma, R. (2025). Context-Sensitive Communication Frameworks for Intelligent Learning Systems. *Transactions on Secure Communication Networks and Protocol Engineering*, 56-64. <https://doi.org/10.17051/TSCNPE/02.02.07>
- Velliangiri, A. (2025). Reinforcement Learning-Based Adaptive Load Forecasting for Decentralized Smart Grids. *National Journal of Intelligent Power Systems and Technology*, 21-28. <https://doi.org/10.17051/NJIPST/01.01.03>
- Verma, S., Tripathi, L., & Sharma, D. (2024, January). Hybrid Models of Machine learning based Multi-Agent Coordination for Optimized Course Selection in E-Learning Environment. In *2024 Fourth International Conference on Advances in Electrical, Computing, Communication and Sustainable Technologies (ICAECT)* (pp. 1-7). IEEE. <https://doi.org/10.1109/ICAECT60202.2024.10469494>