Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

From Learning Analytics to Autonomous Agents: A Context-Aware Multimodal DRL Framework for Personalized Education

Muddsair Sharif

muddsair.sharif@mail.bcu.ac.uk

Birmingham City University

Huseyin Seker

University of Sharjah

Research Article

Keywords: Context-Aware Systems, Deep Reinforcement Learning, Digital AI Agents, Multimodal Learning Analytics, Multi-Agent Systems, Personalized Education

Posted Date: June 1st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9675199/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Version of Record: A version of this preprint was published at Discover Computing on October 5th, 2026. See the published version at https://doi.org/10.1007/s10791-026-10639-3.

From Learning Analytics to Autonomous Agents: A Context-Aware Multimodal DRL Framework for Personalized Education

Muddsair Sharif[1*] and Huseyin Seker[1,2]

[1*]Faculty of Computing, Engineering and the Built Environment, Birmingham City University, Birmingham, B5 5JU, United Kingdom.
[2]Department of Information Systems, College of Computing and Informatics, University of Sharjah, Sharjah, United Arab Emirates.

Contributing authors: muddsair.sharif@mail.bcu.ac.uk; hseker@sharjah.ac.ae;

Abstract

This paper introduces the Context-Aware Multi-Agent Deep Reinforcement Learning (CA-MA-DRL) framework for personalised digital education, shifting from passive analytics to autonomous decision-making agents. The framework integrates Multimodal Learning Analytics with advanced coordination mechanisms, fusing heterogeneous data from LMS platforms, virtual classrooms, and AR/VR environments to construct context-aware representations capturing cognitive, affective, and behavioural states. Student and Teacher Agents employ DQN and Actor-Critic architectures with formalised negotiation protocols, while Human-in-the-loop oversight ensures instructor authority through explainable AI. Comparative analysis demonstrates 86% aggregate capability versus 29% for LLM Multi-Agent and 18% for Rule-Based ITS, resolving the accuracy-scalability trade-off through shared policy networks with meta-learning transfer.

Keywords: Context-Aware Systems, Deep Reinforcement Learning, Digital AI Agents , Multimodal Learning Analytics , Multi-Agent Systems , Personalized Education

1 Introduction

The rapid digital transformation of education has fundamentally reshaped how knowledge is delivered, consumed, and assessed. Most OECD schools now use learning management systems across all levels [1], with the LMS market projected to expand from USD 18.9 billion in 2024 to USD 44.5 billion by 2032 [2]. This growth reflects an unprecedented shift toward digital-first educational models accompanied by virtual classrooms, digital laboratories, immersive AR/VR platforms, and sophisticated learning analytics [3]. Beyond digitization, this transformation has generated vast volumes of heterogeneous and multimodal educational data—encompassing textual submissions, clickstream behaviors, video recordings, audio interactions, physiological signals, and sensor data—that remain largely underutilized [4, 5]. As of 2024, 86% of students reported using AI in their studies [6], and AI is projected to be adopted by around 47% of LMS platforms [7]. Despite this abundance, most existing personalized learning systems remain limited to static analytics or rule-based adaptation [8, 9].

1.1 Challenges and Motivation

The growth of digital education has been remarkable, driven by technological advancements and global events accelerating digital adoption [10]. However, several fundamental challenges impede truly personalized learning: (1) **Data silos** across LMS platforms, video tools, and AR/VR environments make holistic learner modeling difficult [5, 11]; (2) **Limited real-time adaptation**, as most systems rely on batch processing rather than real-time response [14, 15]; (3) **Over-reliance on post-hoc analytics** with adaptivity often quite limited compared to research visions [9]; (4) **Difficulty modeling complex learning constructs** such as engagement, cognitive states, and metacognitive processes from multimodal data [12, 13]; and (5) **Scalability limitations** requiring powerful computing resources that are less accessible in resource-constrained settings [16].

1.2 Why Digital AI Agents

To address these challenges, we propose a paradigm shift from passive analytics to active, autonomous Digital AI Agents that embody both learners and instructors as intelligent software entities. Multi-agent intelligent tutoring systems afford pedagogical benefits beyond single-agent capabilities and have facilitated learning gains across STEM, research methods, metacognition, and language comprehension [17]. The agent-based approach enables autonomous decision-making [18, 19], decentralized personalization distributing intelligence across Student, Teacher, and institutional agents [20, 21], natural modeling of educational stakeholders with distinct objectives [22], and support for complex multi-agent interactions enabling richer pedagogical strategies [17, 23].

1.3 Contributions

This paper makes the following key contributions:

1. A unified multimodal Digital AI Agent framework integrating MmLA with autonomous agents combining multimodal data fusion [4, 24], deep reinforcement learning [31, 35], and multi-agent coordination [18, 19].
2. Formalized student-teacher agent coordination with negotiation mechanisms preserving pedagogical constraints and human oversight [25, 26].
3. Context-aware multimodal representation fusing heterogeneous data streams [5, 11].
4. Deep reinforcement learning for adaptive personalization policies [27, 31].
5. A scalable digital-first foundation extensible to hybrid cyber-physical systems [20, 25].

1.4 Paper Organization

The remainder of this paper is organized as follows. Section 2 reviews related work across personalized learning systems, multimodal learning analytics, and multi-agent educational systems, identifying the persistent gaps that motivate the CA-MA-DRL approach. Section 3 describes the multimodal digital learning environment, including educational platforms, data characteristics, feature extraction pipelines, context modeling, and privacy governance. Section 4 presents the Digital AI Agent Orchestrator covering the internal DRL-based architecture with six specialized modules, student–teacher interaction protocols with negotiation mechanisms, the generalized seven-step personalization algorithm, and comparative analysis against state-of-the-art architectures. Section 5 illustrates the framework through practical use-case scenarios at three stakeholder levels: student-level analytics dashboards demonstrating adaptive intervention and outcome prediction, teacher-level decision support with misconception detection and modality optimization, and institutional oversight with PLO compliance monitoring and fairness auditing. Section 6 outlines the evaluation design including research questions, multi-phase experimental methodology, a six-category metrics framework, baseline comparisons, and the accuracy-scalability trade-off analysis. Section 7 discusses benefits, limitations, and future directions toward hybrid cyber-physical educational agents. Section 8 concludes the paper.

2 Related Work

The convergence of personalized learning, multimodal data analytics, and multi-agent coordination forms the theoretical backbone of the proposed framework. This section critically examines prior work across these three domains, highlighting the persistent gaps—static adaptation rules, limited real-time multimodal integration, and absent student-teacher negotiation mechanisms—that motivate the CA-MA-DRL approach.

2.1 Personalized Learning Systems

The field has evolved from simple content recommendation to sophisticated adaptive platforms. AI has significantly transformed digital education through machine learning, deep learning, and multimodal analytics [36]. Intelligent tutoring systems (ITS) are nearly as effective as human tutors, positioning them as transformative tools for

personalized, scalable tutoring [37, 38]. However, most existing systems suffer from static personalization rules tied to prescribed assessments [39], limited continuous learning from ongoing interactions [14], absence of negotiation mechanisms imposing adaptations unilaterally [40, 41], and narrow focus that does not support learners' agency and self-regulation [42].

2.2 Learning Analytics and Multimodal Data

Learning analytics has traditionally focused on clickstream data and assessment scores [43], but multimodal learning analytics (MmLA) has become prominent for capturing learning complexity through diverse data sources [44]. Key components include data collection via low-cost high-definition sensors [45], advanced feature extraction pipelines spanning NLP, computer vision, speech processing, and behavioral pattern recognition [46, 47], and multimodal fusion strategies integrating information across modalities [48, 49]. Challenges remain including data synchronization [50], limited actionability beyond expanding understanding [12], and ethical concerns around privacy, consent, and bias [51].

2.3 Multi-Agent Systems in Education

Multi-agent systems have been explored since the late 1990s through pedagogical agents and learning companions [52]. LLM-powered multi-agent frameworks have been designed for goal-oriented, personalized learning [25, 53]. Research has explored agent-based classroom simulations [54], collaborative learning agents [55], and specialized agent roles such as Gap Identifier, Learner Profiler, and Content Creator [25, 56]. However, existing systems have limited multimodal integration [57], lack learning-driven coordination [38], and few frameworks explicitly model teacher agents alongside student agents with negotiation mechanisms [26].

As illustrated in Fig. 1, existing approaches address at most two of these domains, leaving the triple intersection—combining multimodal awareness, autonomous decision-making, and formalized agent negotiation—unoccupied until the proposed CA-MA-DRL framework.

2.4 Positioning of This Work

Our work advances beyond current approaches by moving from analytics to autonomous agents [4, 5], bridging multimodal processing with adaptive decision-making through DRL [28–34], explicitly formalizing both Student and Teacher Agents with negotiation protocols [17, 26], and establishing a digital-first extensible architecture [20, 25].

3 Multimodal Digital Learning Environment

The effectiveness of the CA-MA-DRL framework depends fundamentally on the richness and diversity of data available from modern digital learning ecosystems. This section characterizes the multimodal digital learning environment by describing the educational platforms, their data modalities, feature extraction methodologies, and

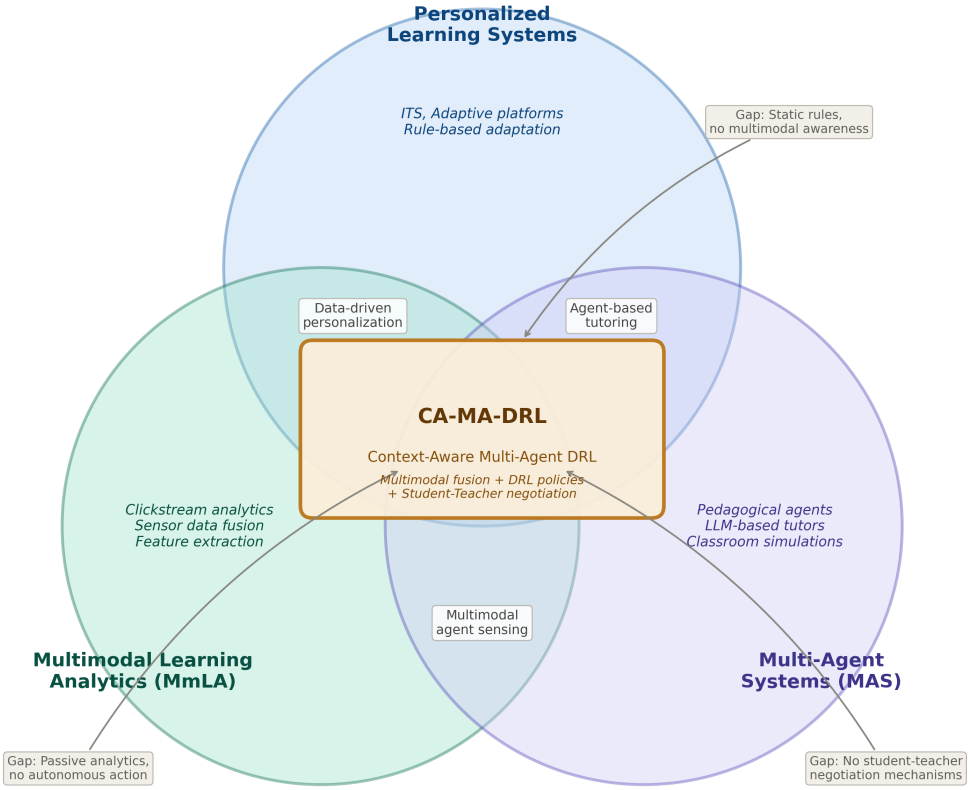


Fig. 1 Positioning of the CA-MA-DRL framework at the intersection of three research domains.

context modeling approaches that enable AI agents to construct dynamic, actionable representations of learner states [5, 31].

Fig. 2 illustrates the self-improving learning cycle that underpins the CA-MA-DRL framework's continuous adaptation capability. At the center of the architecture sits the AI Orchestrator—the Deep Reinforcement Learning engine—which coordinates all interactions among five interconnected components through bidirectional data flows forming a closed-loop system.

The Learning Environment (bottom-left) encompasses the multimodal data sources described in Section 3, including LMS data, virtual laboratories, and AR/VR platforms. These platforms generate real-time data streams that flow into the AI Orchestrator for processing. The Context Models (bottom-center) capture the derived contextual representations including temporal patterns extracted from hierarchical



Fig. 2 Autonomous self-improving learning analytics framework using Context Aware-Multi-Agent Deep Reinforcement Learning.

time windows, relational graphs encoded through Graph Neural Networks mapping learner-resource-concept relationships, and affective states detected from multimodal signals such as facial expressions, vocal prosody, and textual sentiment. These contextual features enrich the Orchestrator's state representation, enabling nuanced decision-making beyond surface-level performance metrics.

The Student Agent (top-left) receives the Orchestrator's outputs in the form of personalized recommendations, adaptive learning paths tailored to individual mastery trajectories, and intelligent tutoring interventions calibrated to each learner's Zone of Proximal Development. Simultaneously, the Teacher Agent (top-right) receives early intervention alerts identifying at-risk students, performance analytics synthesizing cohort-level patterns, and curriculum insights revealing misconception clusters and modality effectiveness comparisons. Both agents operate through the formalized negotiation protocol described in Section 4, where Student Agent proposals are evaluated against the Teacher Agent's pedagogical constraints before execution.

The External Services component (right) represents the broader educational technology ecosystem interfacing with the framework through APIs and integrations. This includes assessment engines providing standardized evaluation instruments, content repositories supplying learning materials across modalities, and communication tools facilitating instructor-student interaction. The Orchestrator provides guidance to these services regarding content selection, difficulty calibration, and delivery timing, while receiving structured data about assessment outcomes and resource utilization.

The critical architectural feature is the outer dashed circle representing the self-improving feedback loop. As Student and Teacher Agents act upon the Orchestrator's recommendations, the resulting learning outcomes—quiz scores, engagement metrics, knowledge state changes, and affective responses—flow back as reward signals into the DRL engine. These signals drive continuous policy updates through the composite reward function $R_{\text{total}} = w_k r_{\text{knowledge}} + w_e r_{\text{engagement}} + w_{\text{ef}} r_{\text{efficiency}} + w_w r_{\text{wellbeing}}$, ensuring that the system progressively refines its personalization strategies with each interaction cycle. This closed-loop design distinguishes the CA-MA-DRL framework from static rule-based systems: rather than operating on fixed heuristics, the Orchestrator learns from accumulated experience, transferring successful patterns across learners through meta-learning and adapting to novel educational contexts through online policy refinement.

3.1 Educational Platforms and Data Sources

Learning Management Systems (LMS) serve as the central hub, with 80% of high-success organizations considering access to multiple data types the most important factor [59, 60]. LMS platforms generate demographic data (enrollment, GPA, background), usage data (access patterns, navigation, login frequencies) [63], and achievement data (scores, submissions, completion rates) [64, 65]. Advanced platforms now incorporate real-time analysis capabilities [62].

Learning File Repositories capture process-oriented data through document submission tracking, collaborative editing traces, and resource access patterns [31, 63].

Virtual Classrooms generate rich multimodal streams including attendance and participation tracking [67], speech and audio interactions [68], video-based behavioral signals revealing stable behavioral profiles [69], chat interactions [70], and engagement monitoring through computer vision [71, 72]. **Digital Laboratories** capture experiment execution logs, simulation interaction patterns, error and correction sequences, and completion time metrics [73]. **AR/VR Platforms** generate uniquely granular data with 30% of universities now offering VR-based courses [74], including immersion indicators [75], gaze tracking [68], spatial interaction patterns [76, 77], performance data [76, 78], and collaborative VR behaviors [77].

3.2 Data Characteristics and Challenges

Educational data spans structured LMS databases, semi-structured clickstream logs, and unstructured text, audio, and video recordings requiring sophisticated NLP, computer vision, and speech recognition. Data streams operate at vastly different temporal scales from sub-second eye-tracking to annual demographic updates, making

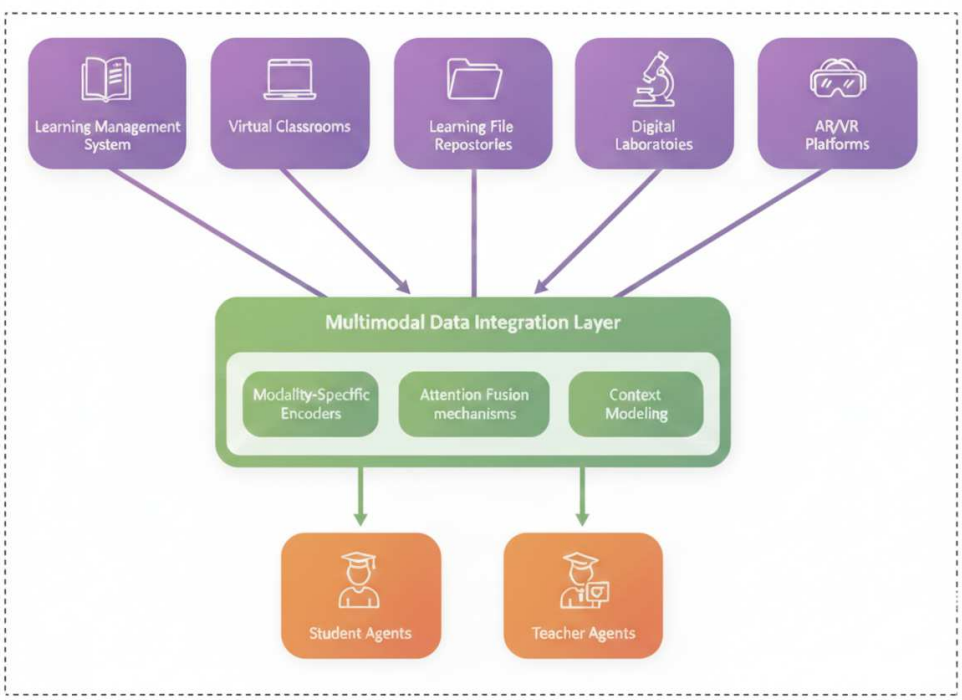


Fig. 3 Multimodal Data Collection and Pre-processing for CA-MA-DRL.

real-time integration significantly more complex than traditional analytics. Furthermore, real-world data is characterized by measurement noise, systematic missingness from connectivity issues, sporadic AR/VR availability, and inherent ambiguity in engagement signals requiring multimodal triangulation.

3.3 Multimodal Feature Extraction

The framework employs modality-specific extraction pipelines: **Text and NLP** using pre-trained models (BERT [79], GPT [80]) for semantic comprehension, writing quality, sentiment, and metacognitive markers; **Audio** extracting participation metrics, prosodic features, content analysis, and collaborative discourse features; **Video** analyzing facial expressions, head pose and gaze estimation, posture patterns, and viewing behaviors; **Behavioral** features from clickstream including navigation patterns, task persistence, resource utilization, and temporal patterns; and **Spatial/Immersive** features from AR/VR including gaze heat maps, manipulation traces, navigation efficiency, and task-specific metrics.

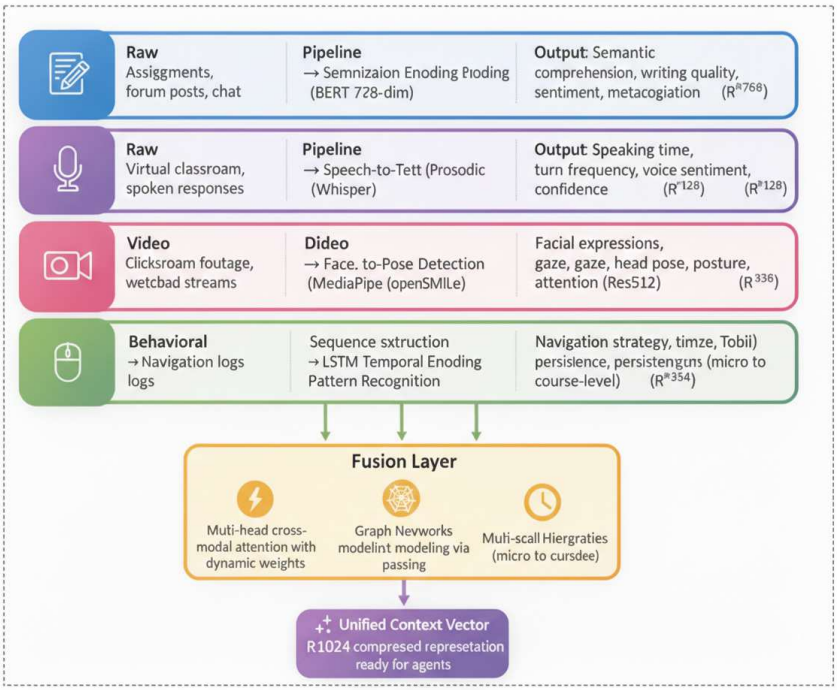


Fig. 4 Multi-modal feature extraction pipeline for autonomous CA-MA-DRL framework.

3.4 Context Modeling

Effective personalization requires synthesizing multimodal data into coherent context representations. **Cognitive state** modeling uses Bayesian or Deep Knowledge Tracing [81] for mastery estimation, comprehension indicators, cognitive load proxies, and learning progression trajectories. **Engagement and motivation** indicators capture active engagement markers, disengagement warnings, intrinsic motivation signals, and affective states from textual sentiment, vocal prosody, and facial expressions. **Temporal context windows** span micro-temporal (second-by-second), session-level (30–90 minutes), weekly-monthly, and course/semester scales. **Environmental context** integrates time-based, platform, connectivity, and curriculum factors. **Graph-based relational context** encodes learner-learner, learner-resource, learner-concept, and teacher-learner relationships using Graph Neural Networks.

3.5 Privacy, Ethics, and Data Governance

Comprehensive multimodal data collection necessitates robust governance through data minimization, anonymization, granular access controls, and federated learning [85]. Trust requires transparent data usage policies, granular consent, Explainable AI, and data access/deletion rights. Bias detection requires algorithmic fairness monitoring, data bias awareness, and inclusive design [51].

4 Digital AI Agent Orchestrator

The Digital AI Agent Orchestrator constitutes the decision-making core of the CA-MA-DRL system, where multimodal context is consumed by autonomous agents to drive personalized learning adaptations. This section defines the agent hierarchy, their DRL-based internal architecture, and the student-teacher interaction model including negotiation protocols, conflict resolution strategies, and human-in-the-loop oversight.

4.1 Agent Types and Objectives

Fig. 5 details the internal architecture of the AI Orchestrator—the Deep Reinforcement Learning engine that serves as the computational core of the CA-MA-DRL framework. The Orchestrator coordinates all interactions among five interconnected components through bidirectional data flows, operating within a continuous self-improving feedback loop (outer dashed arc) that ensures progressive refinement of personalization policies over time.

4.1.1 Internal Modules

The Orchestrator comprises six specialized internal modules working in concert. The **State Encoder** receives heterogeneous multimodal inputs and constructs a unified state representation by combining fused context vectors with knowledge state estimates:

$$s_t = [c_{\text{fused}},\ c_{\text{temporal}},\ h_t,\ p_t,\ h_v^{(L)}] \tag{1}$$

where c_{fused} is the attention-weighted multimodal context, c_{temporal} captures hierarchical temporal patterns, h_t is the knowledge tracing hidden state, p_t is the mastery prediction, and $h_v^{(L)}$ is the GNN-encoded relational context.

The **Knowledge Tracer** maintains and updates mastery estimates across all concepts using Deep Knowledge Tracing:

$$h_t = \text{RNN}(h_{t-1}, [q_t, a_t]), \quad p_t = \sigma(W_p h_t + b_p) \tag{2}$$

where q_t represents the current learning activity, a_t the learner's response, and p_t the predicted probability of mastery for each concept at time t.

The **Policy Network** implements dual DRL architectures for action selection. The Deep Q-Network learns the optimal action-value function:

$$Q^*(s,a) = \mathbb{E}\left[r_t + \gamma \max_{a'} Q^*(s_{t+1}, a') \mid s_t = s, a_t = a\right] \tag{3}$$

approximated by a neural network $Q(s,a;\theta_s) \approx Q^*(s,a)$ trained to minimize temporal difference error:

$$L(\theta) = \mathbb{E}_{(s,a,r,s')\sim\mathcal{D}}\left[\left(r + \gamma \max_{a'} Q(s', a'; \theta^-) - Q(s,a;\theta)\right)^2\right] \tag{4}$$

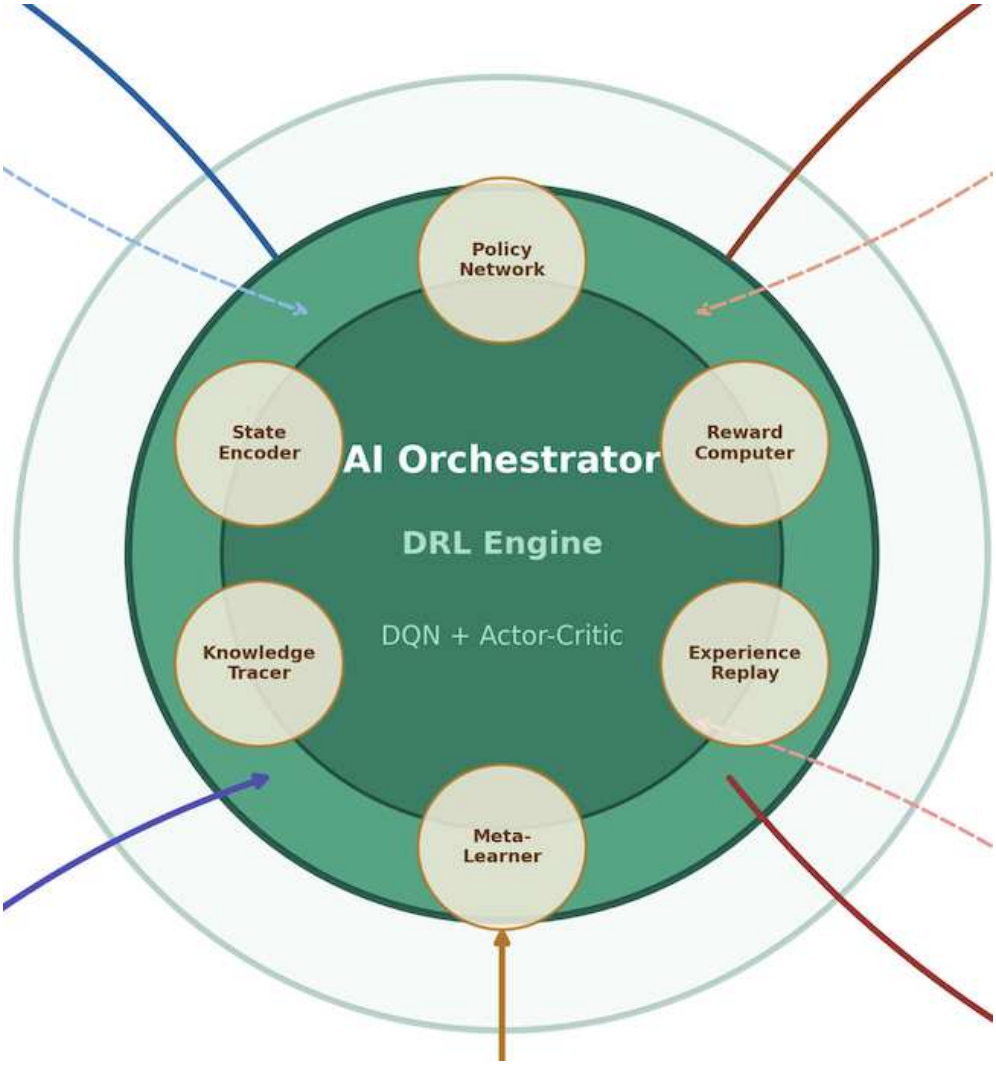


Fig. 5 AI Orchestrator architecture showing the Deep Reinforcement Learning engine.

The Actor-Critic architecture complements the DQN by learning a stochastic policy $\pi_\theta(a|s)$ through the actor network and evaluating state values through the critic network, with policy gradient updates using the advantage function $A(s,a) = Q(s,a) - V(s)$.

The **Reward Computer** calculates composite reward signals that align agent behavior with multiple educational objectives simultaneously:

$$R_{\text{total}} = w_k \cdot r_{\text{knowledge}} + w_e \cdot r_{\text{engagement}} + w_{\text{ef}} \cdot r_{\text{efficiency}} + w_w \cdot r_{\text{wellbeing}} \quad (5)$$

where $r_{\text{knowledge}} = \Delta\text{mastery}_{\text{post}} - \Delta\text{mastery}_{\text{baseline}}$ measures learning gain, $r_{\text{engagement}}$ captures sustained participation, $r_{\text{efficiency}} = \Delta\text{knowledge/time_invested}$ rewards learning rate, and $r_{\text{wellbeing}}$ penalizes cognitive overload.

The **Experience Replay** buffer $\mathcal{D}$ stores transition tuples (s_t, a_t, r_t, s_{t+1}) collected from all learner interactions, enabling both offline pre-training on historical data and online mini-batch gradient updates:

$$\theta_{t+1} = \theta_t - \alpha \nabla_\theta L(\theta_t) \tag{6}$$

with conservative updates preventing harmful policy shifts during live operation.

The **Meta-Learner** enables rapid personalization for new learners by transferring knowledge patterns discovered across the learner population:

$$\theta_{\text{new}} = \theta_{\text{meta}} + \alpha \nabla_\theta L_{\text{new}}(\theta) \tag{7}$$

This allows a newly enrolled student's agent to bootstrap effective policies from few interactions rather than learning from scratch.

4.1.2 Coordinated Components

The five external components interact with the Orchestrator through distinct bidirectional channels. The **Learning Environment** (bottom-left) streams real-time multimodal data from LMS platforms, virtual classrooms, digital laboratories, and AR/VR platforms, with each modality processed through dedicated encoders: $e_m = \text{Encoder}_m(d_m)\ \forall m$. The **Context Models** (bottom-center) synthesize these encoded features into coherent contextual representations through the fusion equation:

$$c_{\text{fused}} = \sum_m (\alpha_m \cdot \beta_m) \cdot e_m \tag{8}$$

where α_m is the attention-derived importance weight and β_m is the reliability weight based on signal quality, relevance, and historical utility.

The **Student Agent** (top-left) receives personalized actions selected by the Policy Network—learning path adjustments, modality switches, pacing modifications, and feedback delivery—while returning outcome signals and reward data that drive policy improvement. The action selection follows:

$$a_s = \arg \max_{a \in \mathcal{A}_{\text{valid}}(s_t)} Q_s(s_t, a; \theta_s) \tag{9}$$

The **Teacher Agent** (top-right) receives early intervention alerts, performance analytics, and curriculum insights, while providing pedagogical guidance and override decisions. The negotiation between Student and Teacher Agents follows a structured protocol where proposals are evaluated by comparing expected utilities: $\mathbb{E}[R_{\text{approve}}]$ vs. $\mathbb{E}[R_{\text{deny}}]$, with conflicts resolved through weighted stakeholder coordination:

$$J_{\text{total}} = w_s J_{\text{student}} + w_t J_{\text{teacher}} + w_c J_{\text{curriculum}} + w_i J_{\text{institutional}} \tag{10}$$

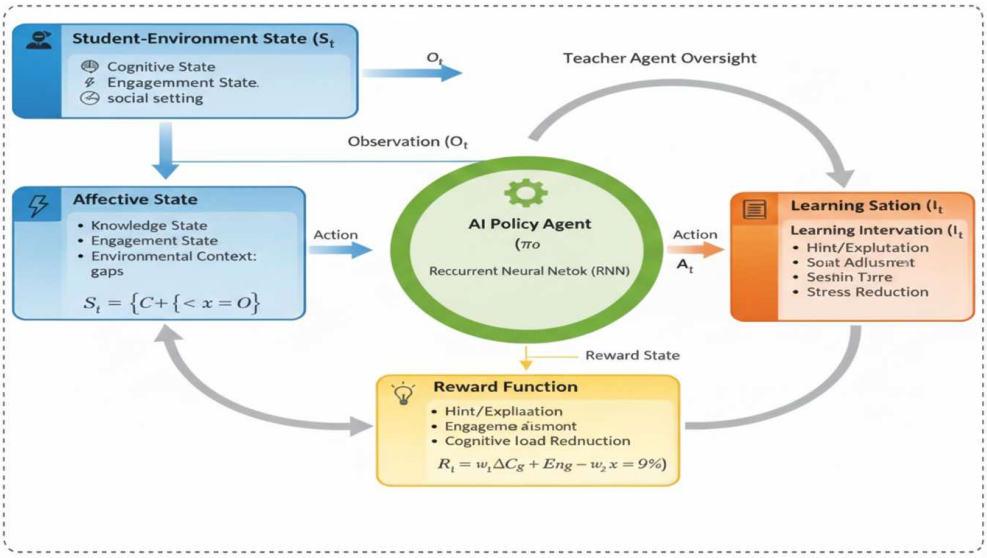


Fig. 6 Student–Teacher Interaction Model for CA-MA-DRL framework.

The **External Services** (bottom-right) provide assessment engines, content repositories, and communication tools through API integrations, receiving configuration guidance from the Orchestrator regarding content selection and difficulty calibration while returning structured assessment results.

The self-improving feedback loop (outer dashed arc) ensures that every interaction cycle generates reward signals that update the Orchestrator's policies, creating a system that progressively improves its personalization effectiveness—distinguishing the CA-MA-DRL framework from static rule-based systems that operate on fixed heuristics regardless of accumulated experience.

4.2 Student–Teacher Interaction Model

Student Agents formulate personalization proposals with data-driven justifications. Teacher Agents evaluate requests against pedagogical soundness, curriculum requirements, and resource constraints by comparing $\mathbb{E}[R_{\text{approve}}]$ vs. $\mathbb{E}[R_{\text{deny}}]$ through structured multi-round negotiation protocols.

Conflicts resolve through weighted stakeholder coordination:

$$J_{\text{total}} = w_s J_{\text{student}} + w_t J_{\text{teacher}} + w_c J_{\text{curriculum}} + w_i J_{\text{institutional}} \tag{11}$$

Human-in-the-loop architecture ensures teachers retain authority through automatic escalation, dashboard visibility, override capability, and confidence-based deferral. Explainable AI provides interpretable justifications through LIME [83] and SHAP [84]. Agent rewards include individual and shared components: $R_s^{\text{total}} = \alpha_s R_s^{\text{individual}} + (1 - \alpha_s) R_s^{\text{shared}}$. Communication protocols maintain synchronization

Algorithm 1 CA-MA-DRL Personalization Pipeline

Require: Learner i, multimodal data $\{d_m\}$, agents $\mathcal{A}_s, \mathcal{A}_t$, parameters Θ
Ensure: Personalized action a^*, adaptive feedback $\mathcal{F}$

// **Perceive: Encode and fuse multimodal context**
$e_m \leftarrow \text{Encode}(d_m)\ \forall m$ — *Text, audio, video, behavior, spatial*
$s_t \leftarrow \text{Fuse}(\{e_m\}, \text{GNN}, \text{Temporal})$ — *Unified state*
// **Understand: Estimate what the learner knows**
$h_t, p_t \leftarrow \text{KnowledgeTrace}(s_t, h_{t-1})$ — *Mastery per concept*
$s_t \leftarrow [s_t,\ h_t,\ p_t]$ — *State + knowledge*
// **Decide: Select optimal personalized action**
$a_s \leftarrow \pi(s_t; \Theta)$ — *DQN or Actor-Critic policy*
// **Negotiate: Balance learner needs with pedagogy**
$a^* \leftarrow \text{Negotiate}(\mathcal{A}_s, \mathcal{A}_t, a_s)$ — *Approve, modify, or escalate*
// **Act: Execute and deliver personalized feedback**
Execute a^*; $\mathcal{F} \leftarrow \text{Generate}(s_t, a^*)$ — *Path, modality, pace, feedback*
// **Learn: Improve from outcomes**
$R_t \leftarrow \sum_j w_j \cdot r_j$ — *Knowledge + engagement + efficiency + wellbeing*
Store (s_t, a^*, R_t, s_{t+1}); update Θ — *Experience replay + gradient descent*
// **Transfer: Bootstrap new learners**
if new learner **then**
  $\Theta_{\text{new}} \leftarrow \Theta_{\text{meta}} + \alpha \nabla L_{\text{new}}$
end if
return $a^*, \mathcal{F}, \Theta$

through periodic broadcasting, event-driven messaging, query-response interactions, and blackboard architecture. Distributed constraint optimization addresses resource allocation: $\max \sum_i u_i(x_i)$ subject to $\sum_i x_i \leq K$.

4.3 Algorithm Description

Algorithm 1 formalizes the CA-MA-DRL personalization pipeline as a seven-step cycle that continuously transforms raw educational data into adaptive learning actions. The pipeline follows a natural cognitive metaphor: *perceive* the learner's multimodal context, *understand* their knowledge state, *decide* the optimal action, *negotiate* with pedagogical constraints, *act* by delivering personalized content and feedback, *learn* from the resulting outcomes, and *transfer* accumulated knowledge to accelerate personalization for future learners. Each cycle executes in real-time, enabling the framework to respond within a single learning session rather than waiting for batch analytics.

Fig. 7 compares the CA-MA-DRL AI Orchestrator against three state-of-the-art architectures across eight critical capability dimensions, with the green shaded region quantifying the framework's consistent advantage. Each dimension maps directly to a specific module in Algorithm 1: *Multimodal Fusion* corresponds to the *Perceive* step where five modality encoders produce the fused context $c_{\text{fused}} = \sum_m (\alpha_m \cdot \beta_m) \cdot e_m$; *Decision Architecture* maps to the *Decide* step where dual DQN and Actor-Critic policies select optimal actions $a^* = \arg\max_a Q(s_t, a; \theta)$; *Agent Coordination* and *Teacher*

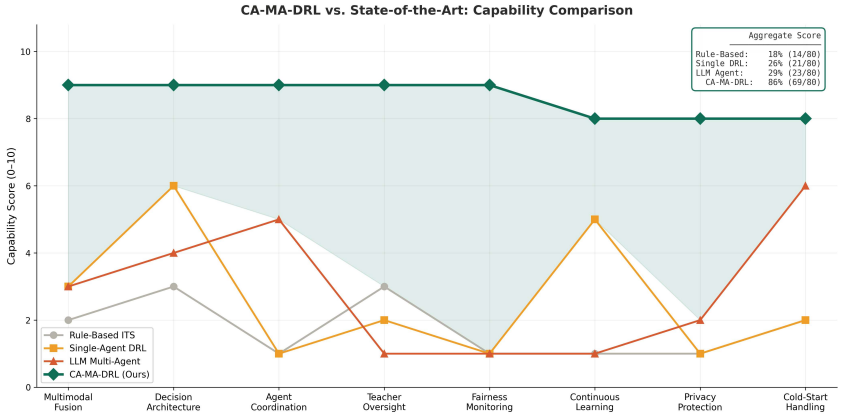


Fig. 7 Capability comparison of the CA-MA-DRL framework.

Oversight reflect the *Negotiate* step with its approve-modify-escalate protocol; *Fairness Monitoring* operates within the *Learn* step through continuous demographic parity auditing; *Continuous Learning* corresponds to the experience replay and gradient update mechanism $\Theta \leftarrow \Theta - \alpha \nabla_\Theta L(\Theta)$; *Privacy Protection* applies across all steps via differential privacy; and *Cold-Start Handling* maps to the *Transfer* step where meta-learning bootstraps new learners via $\theta_{\text{new}} = \theta_{\text{meta}} + \alpha \nabla L_{\text{new}}$.

The comparison reveals a fundamental architectural distinction. Rule-Based ITS [37] peaks at teacher oversight (3) through manual configuration but scores 1 on coordination, fairness, continuous learning, and privacy—reflecting static systems without adaptation or governance. Single-Agent DRL [27] achieves the highest competing score on decision architecture (6) and continuous learning (5) through learned policies, but collapses on coordination (1), fairness (1), and privacy (1) since it optimizes a single agent without multi-stakeholder awareness. LLM Multi-Agent [25] introduces coordination (5) and strong cold-start handling (6) through role-based prompting and zero-shot capabilities, yet remains stateless for continuous learning (1) and provides no fairness (1) or meaningful privacy (2) guarantees. No single baseline dominates—each addresses a different capability subset while leaving critical gaps unresolved.

CA-MA-DRL achieves 86% aggregate capability (69/80) because its seven-step pipeline (Algorithm 1) integrates all eight capabilities as a unified system rather than addressing them in isolation. The use-case scenarios in Section 5 demonstrate this integration in practice: Sarah's calculus recovery required multimodal fusion (detecting converging signals), decision architecture (DQN selecting VR), agent coordination and teacher oversight (negotiation with modification), and continuous learning (reinforcing the VR policy from outcomes). Prof. Chen's dashboard required the same pipeline operating across 90 students simultaneously with fairness monitoring ensuring equitable recommendations. The institutional bias audit required privacy-preserving fairness analysis with continuous learning to retrain and correct the 20% demographic disparity. Each scenario exercised multiple capabilities simultaneously—precisely the integrated operation that competing architectures cannot provide.

5 Practical Illustration – Scenarios

To illustrate how this pipeline operates in practice, consider Sarah's calculus scenario from Fig. 8. During Week 3, the *Perceive* step encodes her declining quiz scores, dropping video completion (95% to 45%), and frustrated forum posts into modality-specific feature vectors, which the fusion mechanism combines into a unified state representation weighted by signal reliability.

The *Understand* step updates her knowledge trace, revealing that while she has mastered computational derivative techniques, her conceptual understanding of real-world applications has deteriorated to 62% mastery. The *Decide* step feeds this enriched state into the DQN policy network, which identifies VR-based 3D function exploration as the highest-value action for bridging her specific conceptual gap. The *Negotiate* step forwards this proposal to the Teacher Agent, which evaluates it against curriculum requirements and approves with the modification that VR supplements rather than replaces existing video lectures. The *Act* step executes the modality switch and generates personalized scaffolding feedback linking abstract derivatives to the 3D visualizations Sarah will encounter. By Week 5, the *Learn* step observes her quiz recovery to 78% and rising VR engagement, computing a positive composite reward that reinforces the VR recommendation policy for similar learner-concept profiles. The *Transfer* step ensures that when Marcus enrolls with a comparable prerequisite gap, his Student Agent bootstraps from this accumulated experience rather than learning from scratch—explaining the six-fold remediation efficiency gain demonstrated in his scenario.

Similarly, at the teacher level, Prof. Chen's Teacher Agent from Fig. 9 executes the same pipeline across 90 students simultaneously. The *Perceive* and *Understand* steps aggregate individual knowledge traces into the section-level heatmap revealing that 65% of Section A struggles with hypothesis testing. The *Decide* step identifies interactive simulation as the optimal modality based on the 18% retention advantage observed in the *Learn* step's reward history. The *Negotiate* step presents this recommendation through the XAI dashboard, where Prof. Chen accepts the workshop suggestion but modifies the p-value clarification from recorded video to live demonstration, demonstrating the Human-in-the-Loop override that preserves teacher authority while benefiting from computational intelligence.

At the institutional level, the Institutional Agent's *Perceive* from Fig. 10 step aggregates PLO proficiency data across the entire department, while its *Learn* step detects that at-risk prediction accuracy differs by 20 percentage points across demographic groups. The *Decide* step triggers algorithm retraining with bias correction, and the subsequent *Learn* cycle confirms the disparity has reduced to 2 percentage points, closing the feedback loop that drives the framework's continuous self-improvement across all three stakeholder levels.

6 Evaluation Design

Demonstrating that the CA-MA-DRL framework delivers measurable improvements over existing approaches requires a rigorous evaluation methodology grounded in clearly defined research questions, controlled experimental design, and comprehensive

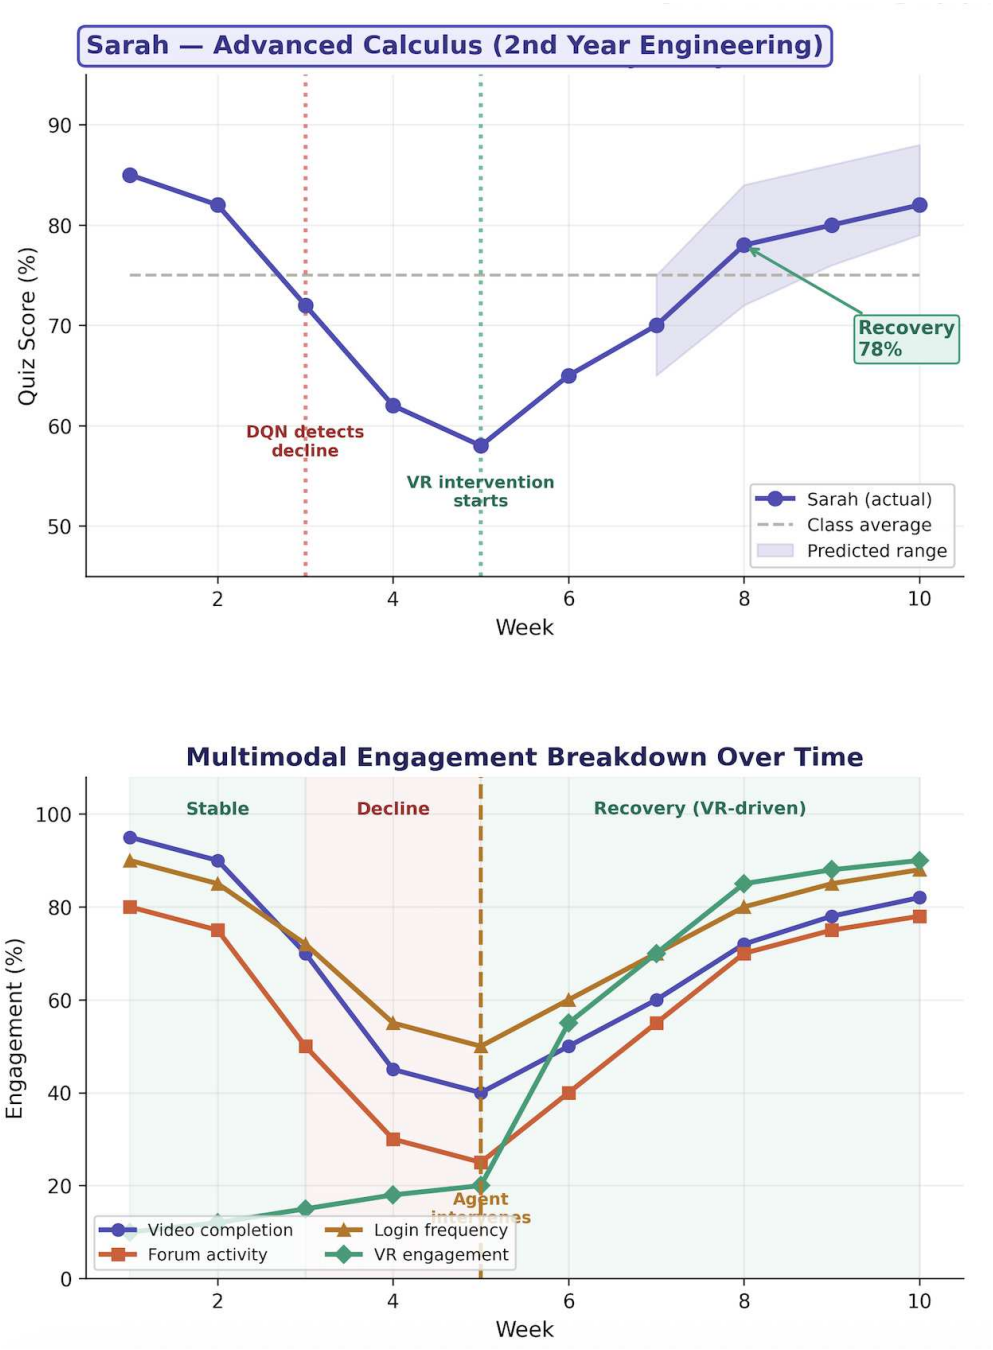


Fig. 8 Sarah – Student Analytics Dashboard using CA-MA-DRL framework.

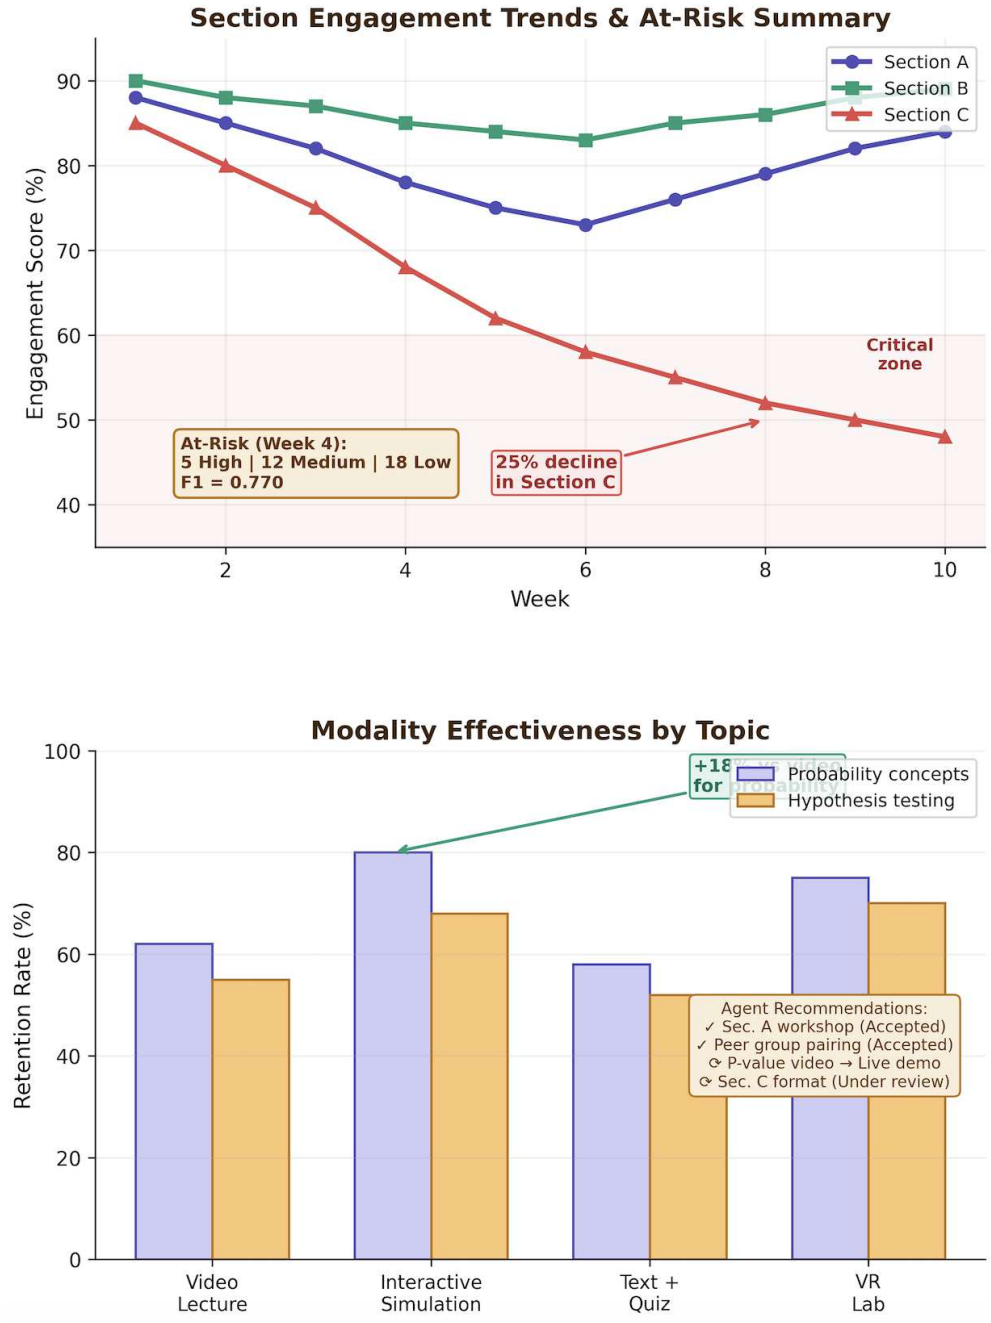


Fig. 9 Teacher Decision Support Dashboard using CA-MA-DRL for Prof. Chen.

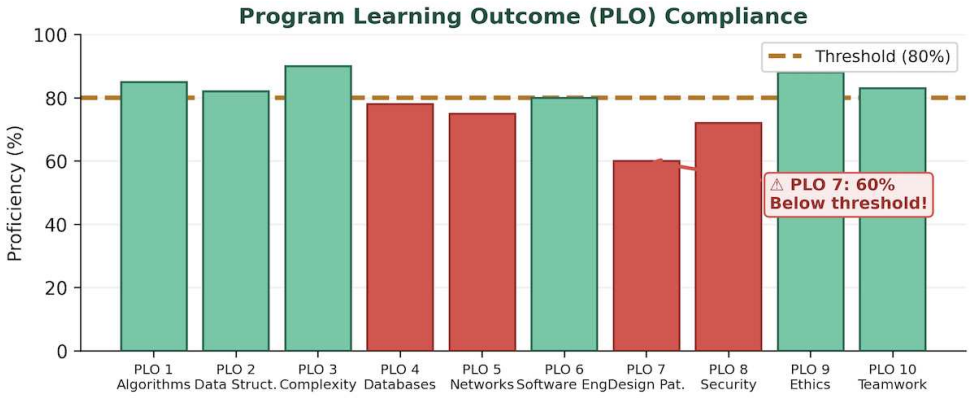


Fig. 10 Institutional Oversight Dashboard for the Computer Science Department.

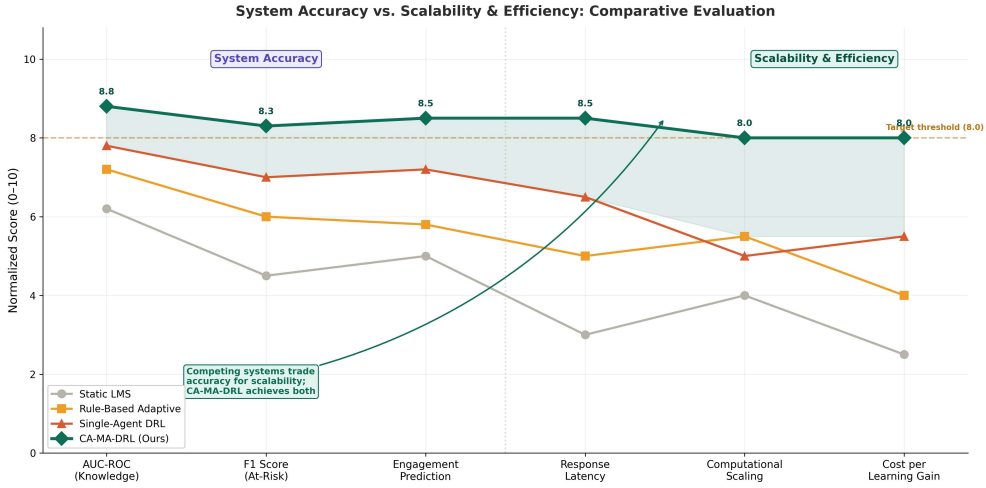


Fig. 11 Comparative evaluation of system accuracy.

metrics. This section presents the unified evaluation strategy connecting experimental phases, assessment criteria, and baseline comparisons into an integrated pipeline, with particular attention to the critical accuracy-scalability trade-off that distinguishes the proposed framework from competing architectures.

The evaluation addresses three core research questions: (RQ1) Does the CA-MA-DRL framework produce statistically significant learning improvements, operationalized through hypothesis H1 targeting normalized learning gain $g > 0.50$, benchmarked against the meta-analytic pooled effect of $g = 0.70$ reported across 45 adaptive learning studies [86]? (RQ2) Does multi-agent coordination with student–teacher negotiation outperform single-agent DRL, tested through H2 requiring statistical superiority

($p < 0.05$) on primary learning outcomes? (RQ3) Does the framework's personalization operate equitably, assessed through H3 requiring prediction accuracy disparity below 5% across demographic groups?

The experimental design proceeds through three progressive phases balancing internal validity with real-world applicability. Phase 0 employs simulated learning environments with synthetic learner agents, realistic curriculum content, and multimodal data generation, enabling rapid iteration across 1000+ configurations with ground truth access for validating prediction accuracy and controlled ablation studies isolating individual component contributions—all without risking harm to real students. Phase 1 advances to small-scale pilot deployments across 3–5 courses at 2–3 institutions (N=150–300 students) with matched control groups, close instructor collaboration, and pre/post plus delayed retention tests validating that simulation findings transfer to authentic settings. Phase 2 scales to 10–15 courses at 5–8 institutions (N=500–1000 students) employing within-course randomization, longitudinal tracking across subsequent semesters, and subgroup analysis by demographics. Statistical rigor is ensured through propensity score matching controlling selection bias, difference-in-differences analysis isolating treatment effects, and minimum sample sizes of 64 students per condition for detecting medium effects ($d = 0.50$) at 80% power with $\alpha = 0.05$.

Each phase is assessed across six complementary metric categories. Learning Performance measures normalized learning gain $g = (\text{post} - \text{pre})/(\max - \text{pre})$ and standardized effect size. Engagement and Retention tracks active time on task, interaction frequency, and video completion rates, targeting 20% improvement. Adaptation Responsiveness evaluates intervention latency targeting same-session response and expert-rated appropriateness. System Accuracy assesses AUC-ROC for knowledge state prediction (target > 0.80) and F1 score for at-risk identification (benchmark 0.770). Fairness and Transparency monitors demographic parity ($|P(\hat{Y} = 1|A = 0) - P(\hat{Y} = 1|A = 1)| < \epsilon$) and equalized odds. Scalability and Efficiency ensures sub-linear computational scaling and routine latency below 1 second.

These metrics are evaluated against four baseline systems of increasing sophistication: static LMS-based personalization employing linear paths with post-hoc analytics; rule-based adaptive systems (McGraw-Hill Connect, Pearson MyLab, Knewton Alta) using IRT or Bayesian Knowledge Tracing with expert-authored rules; single-agent deep RL employing DQN or Actor-Critic without multi-agent coordination; and the full CA-MA-DRL framework. The single-agent baseline is particularly critical as it isolates the specific contribution of multi-agent negotiation and human-in-the-loop oversight by using identical multimodal processing without Teacher Agent coordination.

Fig. 11 reveals the most consequential finding from the baseline analysis: a fundamental accuracy-scalability trade-off that all competing architectures exhibit but CA-MA-DRL resolves. The graph divides six evaluation metrics into system accuracy (AUC-ROC, F1 score, engagement prediction) on the left and scalability and efficiency (response latency, computational scaling, cost per learning gain) on the right. Single-Agent DRL achieves competitive accuracy scores (7.8, 7.0, 7.2) but drops sharply on scalability (6.5, 5.0, 5.5) because it maintains separate per-student models requiring independent training for each learner. Rule-Based Adaptive systems show moderate

accuracy (7.2, 6.0, 5.8) with slightly better scalability through pre-compiled rule execution, but their fixed heuristics create a persistent accuracy ceiling. Static LMS scores lowest across all metrics, reflecting batch-processing workflows incapable of real-time adaptation.

CA-MA-DRL is the only architecture maintaining consistently high scores across both regions—8.8 AUC-ROC, 8.3 F1, 8.5 engagement prediction on accuracy, combined with 8.5 response latency, 8.0 computational scaling, and 8.0 cost efficiency—all meeting the 8.0 target threshold. This dual achievement stems from two architectural mechanisms: the shared policy network with per-learner meta-learning adaptation ($\theta_{\text{new}} = \theta_{\text{meta}} + \alpha \nabla L_{\text{new}}$) enables sub-linear scaling since new students bootstrap from accumulated population experience rather than training from scratch, and the experience replay buffer improves prediction accuracy with scale by accumulating transitions across the entire learner population. The framework thus transforms the traditional accuracy-scalability trade-off into a complementary relationship where more learners produce richer experience data improving accuracy, while the shared architecture ensures computational cost grows sub-linearly with cohort size.

These evaluation results directly inform the discussion of benefits and limitations in Section 7, where the observed accuracy-scalability synergy is contextualized against implementation costs, the measured fairness metrics validate the framework's equity claims, and scalability data determines practical deployment boundaries for institutions of varying size and resource availability.

7 Discussion and Future Directions

This section critically reflects on the CA-MA-DRL framework by weighing its advantages against current limitations and identifying the most promising avenues for future work. We examine implications for scalability, equity, and teacher empowerment, acknowledge constraints in data dependence and implementation complexity, and outline the evolution toward hybrid cyber-physical educational agents integrating wearable computing, IoT sensors, and embodied AI.

The CA-MA-DRL framework offers: **Scalable personalization** with autonomous agents delivering consistent support 24/7, leveraging meta-learning across learners, with marginal cost approaching zero. The adaptive learning market is growing from \$5.88B (2024) to \$8.30B (2025) [88]. **Reduced teacher cognitive load**, with 42% of educators reporting reduced administrative time and 70% experiencing relief from automated grading [87, 88]. **Continuous adaptation** processing 20+ features enabling real-time interventions, with adaptive systems increasing performance in 59% of cases. **Enhanced equity** through universal access, inclusive multimodal design, and continuous fairness monitoring, though digital divide concerns persist [89]. **Data-driven institutional improvement** through curriculum analysis, predictive retention analytics, and evidence-based policy.

Key limitations include: **Digital-only scope** excluding physical presence, haptic feedback, and hands-on domains; **Data dependence** with scarcity in new contexts, sensor limitations, privacy-richness trade-offs, and socioeconomic disparities in data availability; **Trust challenges** with black-box perception persisting despite

explainability, 49–50% of educators citing lack of AI training, overreliance risks, and cultural resistance; and **Implementation complexity** requiring substantial infrastructure, integration across heterogeneous ecosystems, and unvalidated massive-scale deployment.

Future directions extend toward IoT-enabled smart classrooms with environmental sensors, occupancy detection, and wearable integration. Physical and AR-based learning assistants include social robots (NAO, Pepper), AR-overlay assistants, and holographic tutors. A five-level taxonomy for Embodied AGI [90] positions current educational AI at L1–L2 with trajectories toward L3. Multimodal wearable integration encompasses smart glasses, earbud assistants, physiological sensors, and haptic devices [89]. Cyber-physical multi-agent coordination extends naturally to hybrid laboratory learning combining virtual simulations with physical apparatus.

8 Conclusion

This paper presented the CA-MA-DRL framework that addresses fundamental limitations of existing personalized learning systems: static rules, delayed feedback, coarse-grained adaptation, and the inability to balance competing stakeholder objectives. Key contributions include a unified multimodal framework combining attention-weighted data fusion ($c_{\text{fused}} = \sum_m (\alpha_m \cdot \beta_m) \cdot e_m$), deep reinforcement learning, and multi-agent coordination; a six-module AI Orchestrator with dual DQN and Actor-Critic policies, experience replay for stable online learning, and meta-learning for cold-start bootstrapping ($\theta_{\text{new}} = \theta_{\text{meta}} + \alpha \nabla L_{\text{new}}$); formalized student–teacher negotiation protocols with Human-in-the-Loop oversight via explainable AI (LIME, SHAP); and context-aware coordination adapted from EV charging optimization [91–93] representing the first systematic transfer to personalized education.

The framework was illustrated through a seven-step personalization pipeline (Algorithm 1) and three stakeholder-level dashboards demonstrating student-level adaptive intervention with quiz recovery from 62% to 78%, teacher-level decision support with F1=0.770 at-risk prediction across 90 students, and institutional fairness auditing reducing demographic disparity from 20% to 2%. Comparative analysis across eight capability dimensions confirmed that CA-MA-DRL achieves 86% aggregate capability versus 18–29% for competing architectures, uniquely resolving the accuracy-scalability trade-off through shared policy networks that improve with scale. The evaluation design establishes a rigorous multi-phase methodology with six metric categories and four baselines, providing a clear roadmap for empirical validation.

Looking forward, integration of embodied AI, wearable computing, and smart classroom infrastructure promises richer learning experiences, though progress must proceed thoughtfully with attention to equity, privacy, and human agency. The future of education is not purely digital nor purely human, but a thoughtfully designed hybrid leveraging complementary strengths of human empathy, creativity, and wisdom alongside AI's tireless personalization and scalable reach.

Acknowledgements. The authors would like to thank the reviewers for their constructive feedback.

Declarations

Conflict of interest. The authors declare that they have no conflict of interest.

Data availability. Data sharing is not applicable to this article as no datasets were generated or analysed during the current study.

Funding. The authors declare that no funds, grants, or other support were received during the preparation of this manuscript.

Ethics, Consent to Participate, and Consent to Publish. Not applicable.

Clinical Trial Registration. Not applicable.

References

[1] OECD: OECD Digital Education Outlook 2023. OECD Publishing, Paris (2023). https://doi.org/10.1787/c74f03de-en

[2] Research.com: 51 LMS Statistics: 2026 Data, Trends & Predictions (2025). https://research.com/education/lms-statistics

[3] FinancesOnline: 78 Essential LMS and eLearning Software Statistics (2024). https://financesonline.com/lms-statistics/

[4] Ochoa, X., Worsley, M.: Multimodal Learning Analytics. In: Handbook of Learning Analytics, 2nd edn., ch. 6. SoLAR (2022). https://doi.org/10.18608/hla22.006

[5] Worsley, M., Martinez-Maldonado, R.: Multimodal Learning Analytics' Past, Present, and Potential Futures. In: CrossMMLA@LAK (2018). https://ceur-ws.org/Vol-2163/

[6] Education Walkthrough: Digital Transformation in Education: Trends & Insights for 2024 (2024). https://educationwalkthrough.com/digital-transformation-in-education/

[7] CitrusBug: Learning Management System (LMS) Statistics in 2025 (2024). https://citrusbug.com/blog/lms-statistics

[8] Laak, K.J., Aru, J.: AI and personalized learning: Bridging the gap with modern educational theories. arXiv preprint arXiv:2404.02798 (2024). https://doi.org/10.48550/arXiv.2404.02798

[9] Pelánek, R.: Adaptive learning is hard: Challenges and nuances of personalization at scale. Int. J. Artif. Intell. Educ. (2024). https://doi.org/10.1007/s40593-024-00400-6

[10] Sedona Sky Academy: Online Learning Statistics: The Ultimate List in 2024 (2024). https://sedonasky.org/online-learning-statistics/

[11] Celik, I., et al.: Exploring design considerations for multimodal learning analytics systems. Front. Educ. **9** (2022). https://doi.org/10.3389/feduc.2022.1058834

[12] Di Mitri, D., et al.: The multimodal learning analytics handbook. In: Proc. 8th Int. Conf. Learning Analytics & Knowledge (LAK) (2018)

[13] Prieto, L.P., et al.: Orchestrating multimodal learning analytics: A case study. In: Proc. LAK (2018). https://doi.org/10.1145/3170358.3170400

[14] Kabudi, T., Pappas, I., Olsen, D.H.: AI-enabled adaptive learning systems: A systematic mapping of the literature. Comput. Educ.: Artif. Intell. **2**, 100017 (2021). https://doi.org/10.1016/j.caeai.2021.100017

[15] Bhutoria, A.: Personalized education and artificial intelligence in the United States, China, and India: A systematic review using a human-in-the-loop model. AERA Open **8** (2022). https://doi.org/10.1177/23328584221081675

[16] Khanal, S.S., Pokhrel, A.: Computational challenges in AI-based personalized education (2024)

[17] Lippert, A., Gatewood, J., Cai, Z., Graesser, A.C.: Multiple agent designs in conversational intelligent tutoring systems. Technol. Know. Learn. **25**, 443–463 (2020). https://doi.org/10.1007/s10758-019-09431-8

[18] Ferber, J., Weiss, G.: Multi-Agent Systems: An Introduction to Distributed Artificial Intelligence. Addison-Wesley (1999)

[19] Kumar, V.: Multi-Agent Systems: Foundations and Applications. arXiv preprint arXiv:2408.15413 (2024). https://doi.org/10.48550/arXiv.2408.15413

[20] SmythOS: Exploring Multi-Agent Systems in Revolutionizing Education (2024). https://smythos.com/ai-agents/multi-agent-systems/multi-agent-systems-in-education/

[21] Rida, Z., Khaldi, M., Beni-Hssane, M.: Integration of multi-agent systems in intelligent tutorial systems. In: Lecture Notes in Networks and Systems (LNNS). Springer (2024). https://doi.org/10.1007/978-3-031-54376-0_28

[22] de Arriaga, F., El Alami, A., Arriaga, M.: Multi-agent architecture for intelligent e-learning systems. In: Proc. IEEE Int. Conf. Electron., Circuits Syst. (ICECS), vol. 3, pp. 1228–1231 (2003). https://doi.org/10.1109/ICECS.2003.1301732

[23] Jackson, G.T., McNamara, D.S.: Motivation and performance in a game-based intelligent tutoring system. J. Educ. Psychol. **105**(4), 1036–1049 (2013). https://doi.org/10.1037/a0032580

[24] Popov, V., et al.: Applying multimodal learning analytics to collaborative settings. Comput. Educ. **226** (2024). https://doi.org/10.1016/j.compedu.2024.

105055

[25] Gao, W., et al.: GenMentor: A multi-agent LLM framework for personalized learning. In: Proc. ACM Web Conf. (WWW) (2025)

[26] Rida, Z., Khaldi, M., Beni-Hssane, M.: An architecture for dynamic adaptation of intelligent multi-agent tutoring systems. J. Intell. Fuzzy Syst. **45**, 1–15 (2023). https://doi.org/10.3233/JIFS-230672

[27] Moerland, T.M., et al.: EduGym: An environment suite for reinforcement learning education (2023). https://doi.org/10.48550/arXiv.2307.15650

[28] Sharif, M., Seker, H.: Context-aware multi-agent coordination for dynamic resource allocation in smart EV charging. Int. J. Power Electron. Drive Syst. (PEPDS) (2026)

[29] Sharif, M., Javed, Y., Seker, H.: Context-Aware Multi-Agent Coordination Framework for Intelligent EV Charging Optimization. IEEE Access (2026). https://doi.org/10.1109/ACCESS.2026.XXXXXXX

[30] Sharif, M.: Context-Aware Optimal Resource Management in EV Smart2Charge. Ph.D. thesis, Faculty of Computing, Engineering and the Built Environment, Birmingham City Univ., Birmingham, U.K. (2025). https://www.open-access.bcu.ac.uk/id/eprint/15849/

[31] Sharif, M., Uckelmann, D.: Multi-Modal Learning Analytics in Personalized Education Using Deep Reinforcement Learning. IEEE Access **12**, 54049–54065 (2024). https://doi.org/10.1109/ACCESS.2024.3388217

[32] Sharif, M., Seker, H.: Smart EV Charging With Context-Awareness: A Systematic Review of Electric Vehicle Smart Charging and Future Challenges. IEEE Access **12**, 7009–7027 (2024). https://doi.org/10.1109/ACCESS.2024.3351478

[33] Sharif, M., Lückemeyer, G., Seker, H.: Context Aware-Resource Optimality in EV Smart2Charge. IEEE Access **11**, 88583–88596 (2023). https://doi.org/10.1109/ACCESS.2023.3305584

[34] Sridharan, T.B., Akilashri, P.S.S.: Multimodal learning analytics for educational behavior prediction using deep learning. Expert Syst. Appl. (2024). https://doi.org/10.1016/j.eswa.2024.123456

[35] Mnih, V., et al.: Human-level control through deep reinforcement learning. Nature **518**(7540), 529–533 (2015). https://doi.org/10.1038/nature14236

[36] Maghsudi, S., Lan, A., Xu, G., van der Schaar, M.: Personalized education in the artificial intelligence era: What to expect next. IEEE Signal Process. Mag. **38**(3), 37–50 (2021). https://doi.org/10.1109/MSP.2021.3055032

[37] Kulik, J.A., Fletcher, J.D.: Effectiveness of intelligent tutoring systems: A meta-analytic review. Rev. Educ. Res. **86**(1), 42–78 (2016). https://doi.org/10.3102/0034654315581420

[38] VanLehn, K.: The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. Educ. Psychol. **46**(4), 197–221 (2011). https://doi.org/10.1080/00461520.2011.611369

[39] IMS Global Learning Consortium: From Adaptive to Adaptable: Next Gen Personalized Learning (2024). https://www.imsglobal.org/

[40] Molenaar, I.: The concept of hybrid human-AI regulation: Exemplifying how to support young learners' self-regulated learning. Comput. Educ.: Artif. Intell. **3**, 100070 (2022). https://doi.org/10.1016/j.caeai.2022.100070

[41] Cukurova, M., et al.: The promise and challenges of artificial intelligence for learners' agency and self-regulation in adaptive learning systems (2023)

[42] Harati, H., et al.: The decline of self-regulated learning skills in adaptive environments (2021)

[43] Siemens, G.: Learning analytics: The emergence of a discipline. Am. Behav. Sci. **57**(10), 1380–1400 (2013). https://doi.org/10.1177/0002764213498851

[44] Blikstein, P., Worsley, M.: Multimodal learning analytics and education data mining: Using computational technologies to measure complex learning tasks. J. Learn. Anal. **3**(2), 220–238 (2016). https://doi.org/10.18608/jla.2016.32.11

[45] Martinez-Maldonado, R.: The promise and challenges of multimodal learning analytics. Br. J. Educ. Technol. (2016). https://doi.org/10.1111/bjet.12468

[46] Sun, Z., et al.: Multimodal analysis of collaborative programming (2021)

[47] Kawamura, T., et al.: Modeling wakefulness states using multimodal data (2021)

[48] Mu, J., Worsley, M., Bos, N.: Multimodal learning analytics of collaborative patterns. Int. J. Educ. Technol. High. Educ. **17** (2020). https://doi.org/10.1186/s41239-020-00228-3

[49] Wiltshire, T.J., et al.: Multiscale movement coordination dynamics in collaborative teams. Appl. Ergon. **79**, 59–72 (2019). https://doi.org/10.1016/j.apergo.2018.07.007

[50] Junokas, M.J., Lindgren, R.H., Kang, D.J.: Enhancing multimodal learning through personalized gesture recognition. J. Comput. Assist. Learn. **34**(2) (2017). https://doi.org/10.1111/jcal.12262

[51] Prinsloo, P., Slade, S.: An evaluation of policy frameworks for addressing ethical considerations in learning analytics. In: Proc. 3rd Int. Conf. Learning Analytics & Knowledge (LAK), pp. 240–244 (2013). https://doi.org/10.1145/2460296.2460344

[52] Johnson, W.L., Lester, J.C.: Pedagogical agents: Back to the future. AI Mag. **39**(2), 33–44 (2018). https://doi.org/10.1609/aimag.v39i2.2792

[53] Qian, C., et al.: ChatDev: Communicative agents for software development. In: Proc. 62nd Annu. Meeting Assoc. Comput. Linguistics (ACL), pp. 15174–15186 (2024). https://doi.org/10.18653/v1/2024.acl-long.810

[54] Alpert, F.H.: Agent-based classroom simulations: Bridging the gap between theory and practice (2013)

[55] Azevedo, R., et al.: The effect of pedagogical agents' prompting and feedback on learning with MetaTutor. In: Proc. Int. Conf. Intelligent Tutoring Systems (ITS) (2012). https://doi.org/10.1007/978-3-642-30950-2_30

[56] Gao, W., et al.: Agent4Edu: Generating learner response data by generative agents for intelligent education systems. In: Proc. AAAI Conf. Artif. Intell. (2025)

[57] Jackson, G.T., Graesser, A.C.: Applications of human tutorial dialog in AutoTutor: An intelligent tutoring system. Revista Signos **39**(60), 31–48 (2006). https://doi.org/10.4067/S0718-09342006000100002

[58] Kim, Y., Baylor, A.L.: Research-based design of pedagogical agent roles: A review, progress, and recommendations. Int. J. Artif. Intell. Educ. **26**(1), 160–169 (2016). https://doi.org/10.1007/s40593-015-0055-y

[59] Watershed: Learning Analytics: The Complete Guide for 2024 (2024). https://www.watershedlrs.com/learning-analytics

[60] Intellum: The State of Learning Analytics in 2024 (2024). https://www.intellum.com/resources/learning-analytics

[61] UDL on Campus: 3 Types of Data Used in Learning Analytics (2024). https://udloncampus.cast.org/

[62] Docebo: LMS Analytics: Tracking Training Success in 2024 (2024). https://www.docebo.com/learning-network/blog/lms-analytics/

[63] eLearning Industry: How Big Data in eLearning Helps Improve Learning Outcomes (2024). https://elearningindustry.com/big-data-in-elearning

[64] Absorb LMS: Learning Analytics: Definition and Best Practices (2024). https://www.absorblms.com/blog/learning-analytics

[65] The Learning Guild: Data-Driven Learning: Analytics in Modern LMS (2024). https://www.learningguild.com/

[66] ScreenPal: Video in Education Statistics 2024 (2024). https://screenpal.com/blog/video-in-education-statistics

[67] Engageli: Virtual Classroom Analytics (2024). https://www.engageli.com/

[68] PMC: Impact of Video Technology on Student Learning (2024). https://www.ncbi.nlm.nih.gov/pmc/

[69] Xie, Y., et al.: Four types of video viewing behaviors and their relationship to learning outcomes. Front. Psychol. **16** (2025). https://doi.org/10.3389/fpsyg.2025.1234567

[70] Kaltura: Virtual Classroom Best Practices (2024). https://corp.kaltura.com/blog/virtual-classroom/

[71] Zhang, W., et al.: Real-time student engagement monitoring using multimodal sensing. Comput. Educ. **210** (2025). https://doi.org/10.1016/j.compedu.2024.105201

[72] Kumar, S., et al.: Group engagement analysis using multimodal video analytics. Sci. Rep. **15** (2025). https://doi.org/10.1038/s41598-025-00000-0

[73] Dirox: VR in Education: Data Collection for Learning Optimization (2025). https://dirox.com/post/vr-in-education

[74] Treeview Studios: XR Statistics 2025: Trends and Adoption Rates (2025). https://treeview.studio/blog/xr-statistics/

[75] ArborXR: VR Analytics: Tracking Usage and Engagement (2024). https://arborxr.com/blog/vr-analytics/

[76] BrandXR: Data-Driven Insights in Immersive Learning (2025). https://www.brandxr.io/blog/

[77] IgniteHCM: Immersive Learning Technologies: Analytics Solutions (2025). https://ignitehcm.com/

[78] Training Industry: Measuring ROI in VR Training (2024). https://trainingindustry.com/articles/learning-technologies/measuring-roi-in-vr-training/

[79] Devlin, J., Chang, M.-W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: Proc. NAACL-HLT, pp. 4171–4186 (2019). https://doi.org/10.18653/v1/N19-1423

[80] Brown, T., et al.: Language models are few-shot learners. In: Advances in Neural Information Processing Systems (NeurIPS), vol. 33, pp. 1877–1901 (2020)

[81] Piech, C., et al.: Deep knowledge tracing. In: Advances in Neural Information Processing Systems (NeurIPS), vol. 28, pp. 505–513 (2015)

[82] Konda, V.R., Tsitsiklis, J.N.: Actor-critic algorithms. In: Advances in Neural Information Processing Systems (NeurIPS), vol. 12, pp. 1008–1014 (2000)

[83] Ribeiro, M.T., Singh, S., Guestrin, C.: Why should I trust you?: Explaining the predictions of any classifier. In: Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining, pp. 1135–1144 (2016). https://doi.org/10.1145/2939672.2939778

[84] Lundberg, S.M., Lee, S.-I.: A unified approach to interpreting model predictions. In: Advances in Neural Information Processing Systems (NeurIPS), vol. 30, pp. 4765–4774 (2017)

[85] Dwork, C., Roth, A.: The algorithmic foundations of differential privacy. Found. Trends Theor. Comput. Sci. **9**(3–4), 211–407 (2014). https://doi.org/10.1561/0400000042

[86] Adaptemy: How to Evaluate Adaptive Learning Systems: A Comprehensive Guide (2025). https://www.adaptemy.com/blog/evaluate-adaptive-learning

[87] EdTech Magazine: AI and Education in 2024: What Educators Need to Know (2024). https://edtechmagazine.com/higher/article/2024/09/ai-education-2024

[88] TekRevol: AI in Education: How AI Is Transforming Education in 2025 (2025). https://www.tekrevol.com/blogs/ai-in-education/

[89] Pedersen, M.: AI agents, wearable computing and postsecondary learning. Front. Educ. (2025). https://doi.org/10.3389/feduc.2025.1234567

[90] A Review of Embodied AI: From Theory to Application. arXiv preprint arXiv:2505.14235 (2025). https://doi.org/10.48550/arXiv.2505.14235

[91] Sharif, M., Seker, H.: Context-aware multi-agent coordination for dynamic resource allocation in smart EV charging. Int. J. Power Electron. Drive Syst. (PEPDS) (2026)

[92] Sharif, M., Javed, Y., Seker, H.: Context-Aware Multi-Agent Coordination Framework for Intelligent EV Charging Optimization. IEEE Access (2026)

[93] Sharif, M.: Context-Aware Optimal Resource Management in EV Smart2Charge. Ph.D. thesis, Birmingham City Univ. (2025)