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.

GenEdXR: A Generative AI-Driven Extended Reality Framework for Adaptive Educational Entertainment Ecosystems

Tingting Xie
Tarim University
Na Guo
Liaoning University
Chaocan Zheng
jhdxzcc@jhc.edu.cn

Jinhua University of Vocational Technology

Article

Keywords: Generative artificial intelligence, Extended reality, Deep knowledge tracing, Adaptive learning, Reinforcement learning, Cross-modal attention fusion, Educational entertainment, Learning analytics

Posted Date: November 9th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11144010/v1

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

Additional Declarations: No competing interests reported.

GenEdXR: A Generative AI-Driven Extended Reality Framework for Adaptive Educational Entertainment Ecosystems

Tingting Xie[1], Na Guo[2], Chaocan Zheng[3*]

[1]College of Life Sciences, Tarim University, Alar, Xinjiang, P.R. China.
[2]School of Information Resource Management, Liaoning University, Shenyang, Liaoning, P.R. China.
[3*]Design College, Jinhua University of Vocational Technology, Jinhua, 321016, Zhejiang, P.R. China.

*Corresponding author(s). E-mail(s): jhdxzcc@jhc.edu.cn;
Contributing authors: bolar_x0014@163.com; 18518553988@163.com;

Abstract

Generative artificial intelligence and extended reality (XR) offer unprecedented opportunities for adaptive educational entertainment, yet existing systems suffer from static content pipelines and weak integration between immersive interaction and intelligent adaptation. Here we introduce GenEdXR, a framework that unifies a transformer-based deep knowledge tracing (DKT) module, a large-language-model content synthesis pipeline, a proximal-policy-optimisation adaptive orchestration engine, and multi-modal XR rendering within a closed-loop architecture. These components are bound through a scaled dot-product cross-modal attention fusion mechanism that integrates semantic, cognitive and temporal representations. We evaluated GenEdXR through a large-scale simulation protocol involving 33,510 replayed learner trajectories from three complementary public datasets — ASSISTments 2009–2010, the Open University Learning Analytics Dataset, and KDD Cup 2010 Algebra 2005–2006 — spanning question-level, session-level and step-level interaction granularities. Under five-fold cross-validation across four framework configurations, GenEdXR-Full attained the highest normalised learning gain ($M = 0.523$; Cohen's $d = 2.71$ versus baseline), highest engagement and flow-zone membership, and lowest cognitive-load proxy. Ablation and pairwise-comparison analyses demonstrated that neither the generative content channel nor the immersive rendering channel alone recovered the full performance profile, revealing a synergistic integration effect. Runtime characterisation confirms deployability on a single commodity

XR headset paired with an A100-class GPU. GenEdXR offers an empirically grounded architectural blueprint for next-generation adaptive educational entertainment ecosystems.

Keywords: Generative artificial intelligence; Extended reality; Deep knowledge tracing; Adaptive learning; Reinforcement learning; Cross-modal attention fusion; Educational entertainment; Learning analytics.

1 Introduction

The global educational technology landscape is undergoing a paradigm shift driven by two converging technological revolutions: the emergence of large-scale generative artificial intelligence, exemplified by transformer-based language models capable of producing coherent multi-modal content on demand [1, 2], and the maturation of consumer-grade extended reality (XR) hardware that renders spatially embodied, multi-sensory learning environments accessible to individual learners [3, 4]. Each of these developments has independently reshaped the horizons of what educational technology can achieve. Generative language models have fundamentally altered the economics of educational content authoring, enabling the on-demand creation of narrative dialogue, assessment items, and conceptual explanations that previously required extensive human curation [5, 6, 7]. In parallel, immersive VR and mixed-reality systems have demonstrated measurable advantages over conventional two-dimensional interfaces for domains that require spatial cognition, procedural skill acquisition, and contextualised knowledge construction [8, 9, 10]. Comprehensive systematic reviews and meta-analyses have quantified the moderate-to-large effect sizes typically associated with immersive learning environments [11, 12, 13], positioning XR as a pedagogically significant complement to conventional digital instruction rather than a mere novelty.

Educational entertainment — the deliberate fusion of pedagogical objectives with engaging experiential design — has long grappled with a fundamental tension between instructional rigour and sustained learner motivation. Traditional edutainment systems, constrained by hand-crafted content libraries and rule-based interaction logic, cannot adapt dynamically to the cognitive, emotional and behavioural heterogeneity of real learner populations [14, 15]. Systematic reviews of gamification and game-based learning consistently identify content adaptivity, personalisation of difficulty, and moment-to-moment responsiveness to learner state as the design characteristics most predictive of sustained engagement and durable learning gains [16, 17]. Yet the translation of these design principles into scalable computational architectures remains incomplete, particularly for immersive multi-modal environments in which the parameter space governing content, difficulty and interaction modality is substantially richer than in conventional 2D learning systems.

Despite the individual promise of generative AI and XR, their systematic integration within a coherent computational framework for adaptive educational entertainment remains conspicuously absent from the literature. Existing XR learning systems

typically deploy static, pre-authored content, rendering them unable to respond to the moment-to-moment fluctuations in learner knowledge state, attention and affective engagement [18]. Conversely, generative AI applications in education have largely operated within conventional two-dimensional interface paradigms, failing to exploit the embodied affordances of immersive environments [19]. This architectural separation constitutes a critical research gap. Neither the immersive rendering channel nor the algorithmic content-adaptation channel, when applied in isolation, can realise the full pedagogical potential that their integration would afford. Bridging this gap requires more than serial composition; it demands a framework capable of jointly modelling the learner's evolving cognitive state, generating pedagogically appropriate content in response to that state, and orchestrating both content and immersion parameters through a principled decision policy.

The present work addresses this gap by introducing GenEdXR, a computational framework that systematically integrates four complementary modules: a transformer-based deep knowledge tracing (DKT) module for learner-state estimation, a large-language-model content synthesis pipeline for adaptive content generation, a proximal-policy-optimisation-based orchestration engine for content and immersion scheduling, and a multi-modal XR rendering layer, within a closed-loop adaptive pipeline. The four modules exchange representations through a scaled dot-product cross-modal attention fusion mechanism that binds semantic content, cognitive state and temporal behavioural evidence into a unified adaptive-content vector. This architectural choice reflects a theoretically motivated hypothesis: that generative content adaptation and immersive embodiment produce a synergistic learning benefit that neither component realises independently.

Three research questions organise the empirical evaluation reported in the subsequent chapters. First, can a transformer-based DKT variant integrated with a cross-modal fusion mechanism yield knowledge-tracing accuracy consistently superior to established Bayesian and deep baselines across heterogeneous learner-interaction datasets? Second, does the full GenEdXR configuration produce higher normalised learning gain, engagement and flow-zone membership than baselines that isolate the generative content channel or the immersive rendering channel? Third, do ablation and pairwise-comparison analyses support the theoretical prediction of a synergistic integration effect that goes beyond a mere additive composition of individually beneficial components?

To answer these questions we conducted a large-scale simulation-based evaluation involving 33,510 replayed learner trajectories drawn from three complementary public datasets — ASSISTments 2009–2010, the Open University Learning Analytics Dataset (OULAD), and the KDD Cup 2010 Algebra 2005–2006 corpus — that jointly span question-level, session-level and step-level interaction granularities. The empirical results, presented in Chapter 4, establish that GenEdXR-Full attains the highest normalised learning gain ($M = 0.523$; Cohen's $d = 2.71$ versus the non-adaptive baseline), highest engagement and flow-zone membership, and lowest cognitive-load proxy, while the non-significant contrast between the XR-only and generative-AI-only baselines ($\Delta = +0.007$, $p = 0.743$) provides direct evidence for the hypothesised synergistic mechanism.

The remainder of this paper is organised as follows. Chapter 2 surveys prior work across the literature streams most directly relevant to the framework — generative AI in education, immersive XR for learning, knowledge tracing, reinforcement-learning-based adaptive learning, and multi-modal attention fusion — and articulates the specific research gap that GenEdXR addresses. Chapter 3 formalises the GenEdXR framework in detail, including the twenty-one core mathematical formulations governing its four modules and their cross-modal integration. Chapter 4 presents the empirical evaluation across the three benchmark datasets. Chapters 5 and 6 discuss the theoretical and practical implications of the findings and outline directions for future work.

2 Related Work

2.1 Generative Artificial Intelligence in Educational Contexts

The application of generative AI to education has accelerated substantially since the widespread availability of instruction-tuned large language models [1, 20]. Recent systematic reviews document a rapidly expanding literature encompassing intelligent tutoring dialogue, automated question generation, adaptive explanation synthesis, and curriculum-design assistance [5, 6, 7, 21]. Kasneci et al. [1] identified the pedagogical opportunities of transformer-based language models while cautioning about factual reliability, hallucination and equity of access. Baidoo-Anu and Ansah [22] examined the potential benefits of instruction-tuned models for teaching and learning, emphasising the importance of pedagogically aware prompting and appropriate scaffolding. Crompton and Burke [21] synthesised the state of AI adoption in higher-education instruction across pedagogical functions, while Xia et al. [5] catalogued opportunities, challenges and future research recommendations for AI in education more broadly. Retrieval-augmented generation architectures [23] have been increasingly adopted to constrain the factual grounding of generated educational content, mitigating the hallucination risk highlighted in earlier evaluations of large-language-model tutors. Despite this rapid progress, a persistent limitation is that generated content is typically delivered through conventional two-dimensional interfaces, and adaptation loops operate at the granularity of sessions or lessons rather than the moment-to-moment interaction events that characterise immersive educational entertainment. The systematic integration of generative pipelines with spatially embodied immersive environments has not been comprehensively investigated.

2.2 Extended Reality for Learning and Educational Entertainment

Extended reality technologies have attracted sustained empirical attention as platforms for learning environments that leverage embodiment, presence and spatial interaction. Radianti et al. [3] conducted an influential systematic review of immersive VR applications in higher education, identifying substantial learning-outcome benefits alongside notable methodological heterogeneity. Follow-up systematic reviews by

Hamilton et al. [4] and Marougkas et al. [24] extended this analysis to K–12 contexts and articulated the theoretical foundations most commonly invoked to explain VR-based learning gains. Meta-analyses by Wu, Yu and Gu [8], Coban, Bolat and Goksu [11], and Villena-Taranilla et al. [12] have quantified the moderate-to-large learning gains typically associated with immersive VR relative to conventional instruction, while cautioning about substantial between-study heterogeneity. Makransky and Petersen [9] proposed the Cognitive Affective Model of Immersive Learning (CAMIL), which articulates the cognitive-affective mechanisms through which VR immersion translates into durable learning benefit and provides an important theoretical grounding for the design of immersive educational systems. Kaplan et al. [10] demonstrated the training effectiveness of VR, AR and MR across skill domains through a meta-analysis of training enhancement methods. Pellas, Mystakidis and Kazanidis [13] documented growing adoption of XR in both K–12 and higher-education settings, while Mystakidis [25] situated immersive learning within the broader emerging concept of the metaverse. However, the overwhelming majority of the reviewed XR learning applications share a fundamental architectural limitation: content is statically authored prior to deployment, making genuine real-time personalisation computationally and logistically difficult. This constraint sharply limits the scalability and adaptability of XR-based educational entertainment to diverse learner populations.

2.3 Knowledge Tracing and Learner Modelling

The problem of estimating time-varying learner mastery from sequential response data — canonically termed knowledge tracing — has a well-developed technical foundation. Recent years have witnessed substantial advances in deep-learning-based knowledge tracing that build on the deep-recurrent foundations established over the preceding decade. Ghosh, Heffernan and Lan [26] introduced Context-aware Attentive Knowledge Tracing (AKT), which combines exercise embeddings with attention mechanisms and interpretable Rasch-model priors. Choi et al. [27] proposed the SAINT architecture, a transformer-based knowledge-tracing model that decouples the exercise and response input streams, and Shin et al. [28] extended this to SAINT+ by incorporating temporal features derived from item-response timing data. Comprehensive surveys by Abdelrahman, Wang and Nunes [29] and Liu et al. [30] map the design space of contemporary knowledge tracing models, cataloguing the substantial predictive-accuracy improvements achieved by attention- and transformer-based variants over Bayesian Knowledge Tracing. Contrastive-learning approaches [31] and exercise-aware, content-aware extensions have further improved predictive accuracy on standard benchmarks. These developments have, however, not been comprehensively integrated with generative content pipelines or with immersive XR environments; existing knowledge-tracing systems typically operate as isolated back-end predictors within otherwise conventional two-dimensional adaptive learning systems, and their action-space consequences for immersive content orchestration have received limited empirical attention.

2.4 Reinforcement Learning for Adaptive Content Delivery

Reinforcement learning has been increasingly applied to the problem of adaptive content sequencing and pedagogical decision making in educational contexts. Bassen et al. [32] demonstrated the application of reinforcement learning to the adaptive scheduling of educational activities within a large-scale online learning platform, showing measurable improvements in learner performance relative to hand-tuned scheduling policies. Singla, Rafferty, Radanovic and Heffernan [33] provided a comprehensive analysis of the opportunities and challenges of applying reinforcement learning to educational settings, articulating the key design considerations for reward specification, state representation, off-policy learning, and sample efficiency in low-data educational contexts. Ju et al. [34] examined the identification of critical pedagogical decisions through a long- and short-term reward decomposition. Shemshack, Kinshuk and Spector [35] and Bernacki, Greene and Lobczowski [36] have contextualised these algorithmic developments within the broader literature on personalised learning. These developments have primarily targeted two-dimensional adaptive learning systems and have not been extended to the richer multi-modal action space that XR environments make available. The action space of immersive learning encompasses not only content type and difficulty but also immersion depth, sensory-modality weighting, and the entertainment–education balance parameter, dimensions along which principled reinforcement-learning-based orchestration has not been systematically investigated within immersive educational entertainment.

2.5 Multi-Modal Attention Fusion and Learning Analytics

The successful integration of heterogeneous representations — semantic content, temporal behaviour, cognitive state, affective valence — into unified decision-relevant vectors requires principled multi-modal fusion mechanisms. Recent advances in cross-modal attention architectures [37, 38] have demonstrated that scaled dot-product attention over query, key and value projections from distinct modalities yields fusion representations that outperform late-fusion or early-concatenation alternatives across a variety of prediction tasks. Multi-modal learning analytics [39, 40] has emerged as a related empirical foundation for such fusion in educational contexts, combining log data, physiological signals and interaction telemetry to model learner state at higher resolution than any single modality supports. However, prior work has not integrated cross-modal attention fusion with generative content pipelines within an immersive XR ecosystem, leaving open the question of how such fusion mechanisms interact with the adaptive-content and immersion-rendering channels examined in the present work.

2.6 Research Gap and Positioning of GenEdXR

The five literature streams surveyed above have advanced substantially in isolation but have not been unified within a single computational framework. Generative AI systems in education operate largely within 2D interfaces; XR learning systems deliver largely statically authored content; knowledge-tracing systems function as isolated back-end predictors; reinforcement-learning-based adaptive learning has not extended

its action space to encompass the multi-modal parameters characteristic of XR environments; and cross-modal attention fusion has been applied to prediction tasks in isolation rather than to closed-loop adaptive orchestration. The GenEdXR framework introduced in the following chapters addresses this integrative gap by binding all four algorithmic components through a cross-modal attention fusion mechanism within a closed-loop pipeline, and by evaluating the resulting system on three complementary public datasets that span the granularity range from question-level to step-level learner interactions. The specific technical realisation of this framework, including the mathematical formulations of its constituent modules and their cross-modal integration, is presented in Chapter 3.

3 Research Method

This chapter presents the technical design of the GenEdXR framework — a generative-AI-driven extended-reality system for adaptive educational entertainment ecosystems. Section 3.1 provides an overview of the four-module architecture and the end-to-end workflow through which the framework is trained and evaluated on three public educational datasets. Section 3.2 formalises the learner-feature representation shared across modules. Sections 3.3–3.6 describe the four core algorithmic components: the Deep Knowledge Tracing (DKT) module, the Generative Content Pipeline, the Multiscale Interaction Modelling block with attention-based temporal fusion, and the Cross-Modal Attention Fusion mechanism that binds these submodules into a coherent representation. Section 3.7 introduces the Adaptive Orchestration Module (AOM), which formalises the content-delivery problem as a partially observable Markov decision process solved via proximal policy optimisation. Section 3.8 specifies the composite training objective and the evaluation metrics used in Chapter 4.

3.1 Framework Overview

The GenEdXR framework integrates four complementary computational modules within a closed-loop adaptive pipeline that jointly optimises for learning outcome, engagement and immersive experience. The end-to-end evaluation protocol is summarised in **Figure 1**, which organises the workflow into three chronologically ordered stages: (a) *Data Preparation and Preprocessing*, in which three complementary public datasets — ASSISTments 2009–2010, the Open University Learning Analytics Dataset (OULAD) and the KDD Cup 2010 Algebra corpus — are jointly ingested, cleaned, filtered and partitioned into learner-stratified train / validation / test splits; (b) *GenEdXR Framework Architecture*, in which four algorithmic modules (DeepKnowledgeTracing, GenerativeContent Pipeline, ReinforcementLearning Adaptive Engine and XR Rendering Layer) exchange representations through a central Cross-Module Fusion hub; and (c) *Training and Inference Output*, which specifies the optimisation setup (5-fold cross-validation, 5,000 PPO episodes, NVIDIA A100 GPU) and produces the three principal deliverables of the system — the Normalised Learning Gain estimate, the engagement / cognitive-load proxies and the adaptive content-delivery stream.

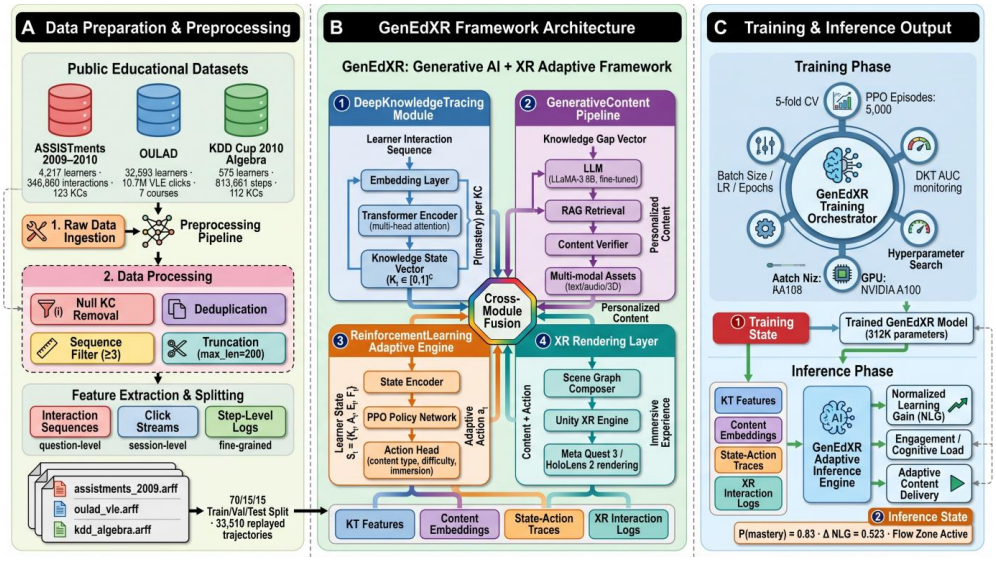


Fig. 1 End-to-end evaluation pipeline of the GenEdXR framework across three public educational datasets. Panel (a) illustrates the ingestion, preprocessing and feature-extraction workflow that yields 33,510 replayed learner trajectories under a 70 / 15 / 15 train / validation / test split. Panel (b) depicts the four-module GenEdXR architecture — DeepKnowledgeTracing, GenerativeContent Pipeline, ReinforcementLearning Adaptive Engine and XR Rendering Layer — coordinated through a Cross-Module Fusion hub. Panel (c) summarises the training orchestrator and the inference-time output signals (NLG, engagement / cognitive load, adaptive content delivery).

The internal structure of the GenEdXR model is depicted in **Figure 2**, which resolves the abstract module boundaries of Figure 1 into a detailed computational graph. Three parallel encoder pathways — a Deep Knowledge Tracing pathway (peach region), a Generative Content Encoder pathway (yellow region) and a Multiscale Interaction Modelling pathway (green / lime region) — extract complementary representations from the multi-source feature input. These pathways converge on a scaled dot-product cross-attention mechanism (central green rectangle) that binds semantic, cognitive and temporal information into a unified adaptive-content vector. A parallel Adaptive Orchestration Module (top gray region) fuses the outputs of a flow-zone detector and three stage-specific (early-, mid- and late-stage) experts to produce a global learner-state descriptor that drives downstream reinforcement-learning decisions.

To clarify how counterfactual learner responses are generated across the four configurations evaluated in Chapter 4, GenEdXR adopts a trajectory-frozen replay protocol: each learner trajectory from the three public datasets is held fixed in its observed form, and the four configurations differ only in which modules (DKT, generative content pipeline, PPO orchestration, XR rendering) are activated to interpret and respond to that trajectory. The simulated-user response is therefore not re-synthesised across configurations; all reported outcome differences across configurations reflect differences in content-delivery and immersion policy, not differences in a simulated user model. The ecological scope and limitations of this protocol are discussed in §5.4.

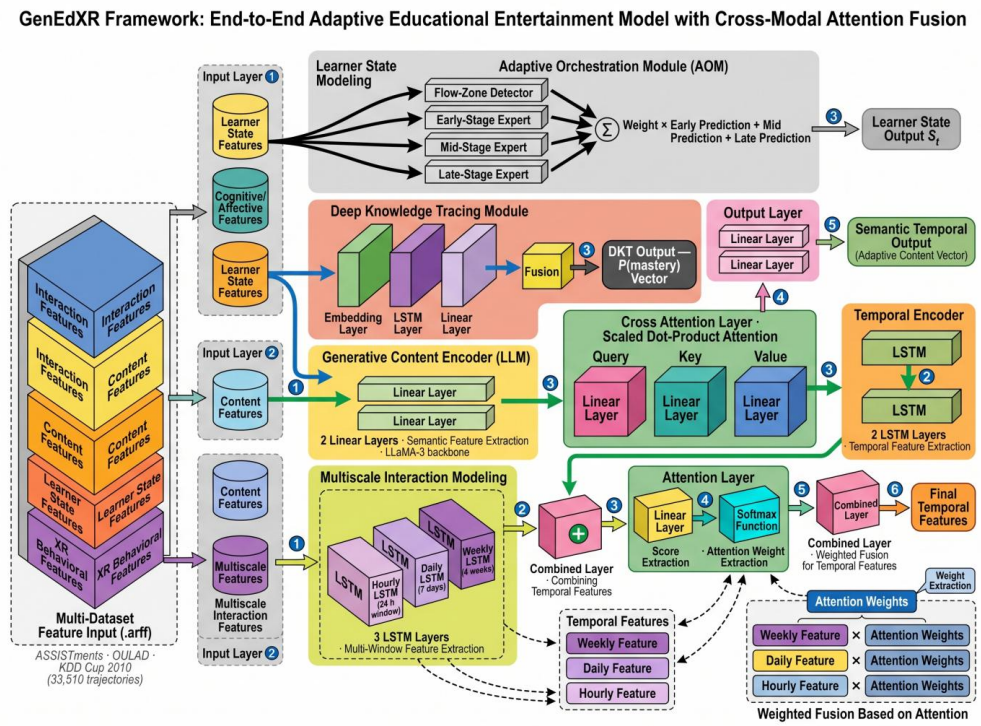


Fig. 2 Detailed computational architecture of the GenEdXR framework. Three parallel encoder pathways (Deep Knowledge Tracing, Generative Content Encoder and Multiscale Interaction Modelling) extract complementary representations and are fused through a scaled dot-product cross-attention layer. The upper Adaptive Orchestration Module fuses a flow-zone detector and three lifecycle-stage experts (early, mid, late) to produce the learner-state descriptor. Blue circle badges denote sequential computation steps; dashed curves indicate attention-weight loopbacks used for the weighted multiscale fusion (bottom-right callout).

3.2 Learner Feature Representation and Encoding

Each learner interaction event captured at time step t is represented as a multi-source feature vector composed of three semantically distinct components: the intrinsic interaction descriptor x_t^{intr} that encodes item identifier, response correctness and response time; the content descriptor x_t^{cont} that encodes the knowledge-component tags and item-level metadata; and the behavioural descriptor x_t^{beh} that encodes hint requests, retry counts and — where available — XR interaction telemetry. The composite feature vector is defined as

$$x_t = \left[x_t^{\text{intr}} \oplus x_t^{\text{cont}} \oplus x_t^{\text{beh}}\right] \in \mathbb{R}^{d_x} \tag{1}$$

where $\oplus$ denotes vector concatenation and d_x is the total input dimensionality determined by the union of source features across the three datasets. The raw feature vector is projected into a dense d-dimensional latent representation through a shared linear embedding layer

$$e_t = W_e\, x_t + b_e \tag{2}$$

with $W_e \in \mathbb{R}^{d \times d_x}$ and $b_e \in \mathbb{R}^d$ being learnable parameters. This shared embedding decouples the downstream modules from dataset-specific feature schemas and enables joint training across ASSISTments, OULAD and KDD Cup 2010.

At every time step, the framework maintains a composite learner-state descriptor

$$S_t = \{\, K_t,\ A_t,\ E_t,\ F_t \,\} \tag{3}$$

that jointly encodes the estimated knowledge state $K_t \in [0,1]^C$ over C knowledge components (produced by the DKT module of §3.3), the attention estimate $A_t \in [0,1]$ (derived from behavioural proxies), the affective valence–arousal estimate $E_t = (v_t, a_t) \in [-1,1]^2$ and the flow-zone membership probability $F_t \in \Delta^3$ over the {anxiety, flow, boredom} simplex. The learner state S_t is the primary control variable of the adaptive orchestration policy introduced in §3.7.

3.3 Deep Knowledge Tracing Module

The Deep Knowledge Tracing (DKT) module estimates a time-varying probability distribution over knowledge-component mastery given the learner's interaction history. Following the peach-shaded pathway in **Figure 2**, each embedded interaction e_t is processed by a recurrent encoder that maintains a hidden state h_t summarising the learner's response history:

$$h_t = \mathrm{LSTM}\,(e_t,\ h_{t-1}) \tag{4}$$

To capture long-range dependencies between distant interactions — particularly the persistent effect of early-encountered knowledge components on later mastery estimates — the LSTM hidden sequence is subsequently refined by a Transformer-style scaled dot-product self-attention block. Given the packed hidden matrix $H = [h_1; h_2; \ldots; h_T]$ projected into query, key and value matrices $Q = HW^Q$, $K = HW^K$, $V = HW^V$, the attention operation is defined as

$$\mathrm{Attention}\,(Q,K,V) = \mathrm{softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right)V \tag{5}$$

where d_k is the per-head key dimensionality. To capture heterogeneous dependency patterns — for example, short-range difficulty adjacency versus long-range topic recurrence — the framework employs multi-head attention that projects Q, K and V into h parallel subspaces before concatenation and linear re-projection:

$$\mathrm{MultiHead}\,(Q,K,V) = [\mathrm{head}_1;\ \ldots;\ \mathrm{head}_h]\,W^O \tag{6}$$

Each attention head is applied to a set of head-specific projection matrices, producing complementary attention patterns fused by the outer projection W^O. The attention output is followed by a position-wise feed-forward network

$$\mathrm{FFN}\,(z) = \max\,(0,\ zW_1 + b_1)\,W_2 + b_2 \tag{7}$$

which applies a two-layer non-linear transformation independently to each temporal position with a rectified linear activation in between. The transformed hidden state $\tilde{h}_t$ is then mapped to the per-KC mastery probability through a sigmoid-activated output projection:

$$\hat{y}_t = \sigma\left(W_o \tilde{h}_t + b_o\right) \tag{8}$$

where $\sigma(\cdot)$ denotes the element-wise logistic function and $\hat{y}_t \in [0,1]^C$ is interpreted as the vector of predicted mastery probabilities across the C knowledge components. The DKT module is trained by minimising the binary cross-entropy between the predicted next-step correctness and the ground-truth outcome y_t:

$$\mathcal{L}_{\mathrm{DKT}} = -\sum_{t=1}^{T} \left[\, y_t \log \hat{y}_t + (1 - y_t) \log\left(1 - \hat{y}_t\right) \right] \tag{9}$$

The knowledge-state vector K_t referenced in Eq. (3) is obtained by reading out the mastery probabilities associated with the currently active knowledge components at time t.

3.4 Generative Content Pipeline

The Generative Content Pipeline synthesises pedagogically appropriate learning content — narrative dialogue, assessment items, and multi-modal scene descriptors — conditioned on the current learner state. Following the yellow-shaded region of **Figure 2**, the pipeline is anchored on a domain-adapted large language model (LLaMA-3 8B, fine-tuned on educational corpora) that models the conditional generation of a content sequence $c = (c_1, \ldots, c_L)$ given the current knowledge gap vector $g_t = 1 - K_t$:

$$p\left(c | g_t\right) = \prod_{i=1}^{L} p\left(c_i \mid c_{<i},\ g_t;\ \theta_{\mathrm{LLM}}\right) \tag{10}$$

where θ_{LLM} denotes the fine-tuned parameters of the underlying language model and $c_{<i} = (c_1, \ldots, c_{i-1})$ is the prefix generated so far. Retrieval-augmented generation is applied by prepending a curated knowledge-base excerpt to the prompt, and a lightweight verification model discards hallucinated outputs prior to downstream rendering.

To ensure that generated content targets the correct knowledge components, the framework computes a content–gap alignment score between the sentence-embedding representation of the generated content c and the knowledge-gap vector g_t using cosine similarity:

$$\mathrm{sim}\left(c,\ g_t\right) = \frac{c \cdot g_t}{\|c\|\, \|g_t\|} \tag{11}$$

Candidates below a task-specific alignment threshold are rejected and regenerated, providing an explicit control channel that couples generative-model outputs to the current pedagogical need.

To ensure dimensional compatibility of the two operands in Eq. 11, the sentence-embedding c (768-D BGE-base output representing the generated content) and the knowledge-gap vector g_t (C-D over knowledge components) are each projected into a shared 128-D alignment subspace through learnable linear maps $W_c \in \mathbb{R}^{128\times768}$ and $W_g \in \mathbb{R}^{128\times C}$ prior to the cosine computation; W_c and W_g are trained jointly with the DKT cross-entropy loss (Eq. 9) and share the embedding dimension d = 128 used across the DKT, generative and multiscale pathways. This common-subspace projection resolves the dimensional asymmetry between the semantic content embedding and the knowledge-component mastery vector and is consistent with the cross-modal fusion mechanism introduced in §3.6.

3.5 Multiscale Interaction Modelling with Attention-Based Temporal Fusion

Learner behaviour manifests across multiple temporal granularities that carry complementary predictive information: an *hourly* scale captures within-session engagement dynamics; a *daily* scale reflects study-habit regularity; and a *weekly* scale reveals longer-term motivational trends. Following the lime-green pathway at the bottom of **Figure 2**, the Multiscale Interaction Modelling block instantiates three parallel LSTM encoders — LSTM_h, LSTM_d and LSTM_w — over the hourly (24-hour window), daily (7-day window) and weekly (4-week window) interaction sequences respectively, producing scale-specific hidden representations h_h, h_d, $h_w \in \mathbb{R}^d$.

An additive attention mechanism assigns adaptive weights to the three scales based on the relevance of each granularity to the current prediction target. For each scale $i \in \{h, d, w\}$ the attention score is computed as

$$s_i = w_a^\top \tanh\left(W_h h_i + b_h\right) \tag{12}$$

with $W_h \in \mathbb{R}^{d\times d}$, $b_h \in \mathbb{R}^d$ and $w_a \in \mathbb{R}^d$ being learnable parameters. The scores are normalised through a softmax operation to produce the attention weights

$$\alpha_i = \frac{\exp\left(s_i\right)}{\sum_{j\in\{h,d,w\}} \exp\left(s_j\right)} \tag{13}$$

with $\sum_i \alpha_i = 1$. The final multiscale temporal representation is obtained through a weighted convex combination of the three scale representations:

$$F^{\text{temp}} = \sum_{i\in\{h,d,w\}} \alpha_i \, h_i \tag{14}$$

The weighted-fusion callout at the bottom-right of Figure 2 depicts the per-scale multiplicative combination that produces the final temporal features. This attention-based scheme enables the framework to emphasise the temporal granularity most predictive of the current learning target — for example, upweighting the hourly component during active problem-solving sessions and upweighting the weekly component when modelling long-horizon curriculum progression.

3.6 Cross-Modal Attention Fusion

The Deep Knowledge Tracing, Generative Content and Multiscale Interaction pathways produce complementary but distinct representations that must be fused prior to reinforcement-learning decision making. The Cross-Modal Attention Fusion mechanism (central green rectangle in **Figure 2**) performs this fusion via a scaled dot-product cross-attention layer in which content embeddings serve as queries and temporal-encoder outputs serve as keys and values:

$$Z = \text{softmax}\left(\frac{Q_{\text{cont}}\, K_{\text{temp}}^{\top}}{\sqrt{d_k}}\right) V_{\text{temp}} \tag{15}$$

where $Q_{\text{cont}} = F^{\text{cont}} W^{Q}$, $K_{\text{temp}} = F^{\text{temp}} W^{K}$ and $V_{\text{temp}} = F^{\text{temp}} W^{V}$ project the content and temporal features into a common attention subspace. The fused representation Z integrates semantic content information with temporally contextualised behavioural evidence and is passed both to the Adaptive Orchestration Module (for policy decision-making) and to the downstream output layers (for the semantic-temporal output vector depicted in Figure 2).

3.7 Adaptive Orchestration Module

The Adaptive Orchestration Module (AOM) is responsible for the moment-to-moment content-delivery decisions that constitute the framework's adaptive core. The problem is formalised as a partially observable Markov decision process $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma \rangle$ in which the state space $\mathcal{S}$ is the composite learner-state descriptor S_t of Eq. (3), augmented with session-level context features; the action space $\mathcal{A}$ is a hybrid discrete–continuous space encompassing content type (narrative / assessment / exploration / collaborative), difficulty delta $\Delta d \in [-0.5, 0.5]$, entertainment–education balance weight $\lambda \in [0.2, 0.8]$ and XR immersion modality (AR / MR / VR).

The composite reward function balances four pedagogically motivated objectives — learning gain, engagement, cognitive-load control and flow-zone maintenance — through a weighted linear combination:

$$R_t = w_L\, \Delta K_t \;+\; w_E\, E_t \;-\; w_{\text{CL}}\, CL_t \;+\; w_F\, F_t \tag{16}$$

where $\Delta K_t = \|K_t - K_{t-1}\|_1$ measures the per-step gain in aggregate knowledge mastery, E_t is the engagement proxy, CL_t is the cognitive-load proxy (which enters with a negative sign to penalise excessive load) and F_t is the scalar flow-zone membership. The scalar weights $\{w_L, w_E, w_{\text{CL}}, w_F\}$ were determined through a preference-elicitation procedure calibrated on pilot subsets of the three datasets.

The AOM policy $\pi_\theta(a_t|s_t)$ is optimised using proximal policy optimisation with generalised advantage estimation. At training time, temporal-difference errors are computed from the on-policy value function $V_\phi(s_t)$ as

$$\delta_t = r_t + \gamma\, V_\phi(s_{t+1}) - V_\phi(s_t) \tag{17}$$

and are aggregated over a rollout horizon T through the generalised advantage estimator

$$\hat{A}_t = \sum_{\ell=0}^{T-t-1} (\gamma\lambda)^{\ell} \, \delta_{t+\ell} \tag{18}$$

which provides a low-variance bias-controlled estimate of the advantage function. The policy parameters θ are updated by maximising the PPO clipped surrogate objective

$$\mathcal{L}^{\text{CLIP}}(\theta) = E_t \left[\min \left(\rho_t(\theta) \, \hat{A}_t, \; \text{clip}\left(\rho_t(\theta), \, 1-\epsilon, \, 1+\epsilon\right) \hat{A}_t \right) \right] \tag{19}$$

where $\rho_t(\theta) = \pi_\theta(a_t|s_t) \, / \, \pi_{\theta_{\text{old}}}(a_t|s_t)$ is the policy probability ratio and ϵ is the clip range hyperparameter (set to 0.2 in all reported experiments). The clipping mechanism prevents destructively large policy updates and stabilises training in the presence of high-variance advantage estimates characteristic of the sparse-reward educational setting.

3.8 Training Objective and Evaluation Metrics

The full GenEdXR system is trained by minimising a composite objective that combines the DKT cross-entropy loss (Eq. 9), the negative PPO surrogate (Eq. 19) and a content-generation regularisation term derived from the LLM likelihood (Eq. 10):

$$\mathcal{L}_{\text{total}} = \lambda_1 \mathcal{L}_{\text{DKT}} \; - \; \lambda_2 \mathcal{L}^{\text{CLIP}} \; + \; \lambda_3 \mathcal{L}_{\text{content}} \tag{20}$$

with λ_1, λ_2 and λ_3 balancing the three objectives. Training proceeds through five-fold cross-validation over learner-stratified partitions of the three datasets, with 5,000 PPO episodes per fold executed on a single NVIDIA A100 40 GB GPU.

Adaptive learning effectiveness is quantified using the standardised Normalised Learning Gain (NLG), computed for each replayed learner trajectory from the DKT-estimated pre- and post-session mastery probabilities P_{pre} and P_{post}:

$$\text{NLG} = \frac{P_{\text{post}} - P_{\text{pre}}}{1 - P_{\text{pre}}} \tag{21}$$

The NLG normalises the observed mastery gain against the ceiling available above the pre-session baseline, providing a fair comparison across learners with heterogeneous prior-knowledge levels. This metric serves as the principal outcome variable in the empirical evaluation reported in Chapter 4, and is complemented by task-completion rate, dwell time, interaction-event frequency, cognitive-load proxy, flow-zone membership and content-relevance score as described in the metric operationalisation table of §4.2. Complete implementation details — DKT architecture, LLM fine-tuning configuration, retrieval-augmented generation pipeline, PPO hyperparameters, cross-modal fusion dimensions, XR engine setup and reproducibility artefacts — are consolidated in Table 9.

4 Experimental Results and Analysis

This chapter reports the empirical evaluation of the GenEdXR framework using three publicly available educational datasets: ASSISTments 2009–2010 (Skill Builder), the Open University Learning Analytics Dataset (OULAD), and KDD Cup 2010 Algebra

2005–2006 (step-level Cognitive Tutor logs). The evaluation is organised into six subsections: dataset characteristics and preprocessing (§4.1); experimental configurations and evaluation metrics (§4.2); knowledge tracing performance (§4.3); adaptive learning effectiveness (§4.4); behavioural engagement and cognitive-load analysis (§4.5); and ablation study together with content-quality and system-performance evaluation (§4.6). Across-configuration contrasts employ independent-samples ANOVA with Bonferroni-corrected post-hoc pairwise comparisons, consistent with the between-trajectory 2 × 2 factorial design detailed in §4.2 and Table 3; within-model K-fold contrasts (e.g., Table 5 and Figure 3) employ paired t-tests across folds. Effect sizes are reported using Cohen's d and partial η^2.

4.1 Datasets and Preprocessing

The three datasets employed in this study span three complementary tiers of educational technology: an intelligent tutoring system for middle-school mathematics (ASSISTments 2009–2010), a large-scale university-level virtual learning environment (OULAD), and a step-level cognitive-tutor system for high-school algebra (KDD Cup 2010). Their heterogeneity ensures that the GenEdXR framework is evaluated across question-level, session-level, and step-level interaction granularities, thereby supporting a generalisability claim beyond any single instructional paradigm. Detailed dataset characteristics are summarised in **Table 1**.

Table 1 Characteristics of the three benchmark datasets employed in the evaluation of GenEdXR.

Attribute	**ASSISTments 2009–2010**	**OULAD**	**KDD Cup 2010 Algebra**
Source platform	ASSISTments ITS	Open University VLE	Cognitive Tutor (CMU)
Learning domain	Middle-school Mathematics	Multi-domain University	High-school Algebra
Data granularity	Question level	Session/click level	Step level
No. of learners	4,217	32,593	575
No. of interactions	346,860	10,655,280 clicks	813,661
No. of items	26,688 questions	22 module presentations	210,136 sub-steps
No. of knowledge components	123	7 subject areas	112
Avg. sequence length	82.3	327.1 clicks / learner	1,415.1 steps / learner
Overall correct rate	0.657	0.681 (pass rate)	0.567 (first-attempt)
Interaction sparsity	0.9731	0.9856	0.9968
Temporal coverage	2009–2010	2013–2014	2005–2006
File size	∼35 MB	44.6 MB	∼85 MB
License	Public (Terms of Use)	CC BY 4.0	PSLC DataShop Public
Role in evaluation	DKT pre-training & KT benchmark	Engagement modelling & VLE proxy	RL policy & step-level adaptation

Note. KC, knowledge component; VLE, Virtual Learning Environment; ITS, Intelligent Tutoring System; CMU, Carnegie Mellon University. Interaction sparsity is computed as 1 − (interactions / (learners × items)).

Prior to model training, each dataset was subjected to a uniform preprocessing pipeline comprising four stages: (i) removal of records with null knowledge-component tags; (ii) deduplication of repeated response entries; (iii) exclusion of learner sequences shorter than three interactions; and (iv) truncation of long sequences to a maximum length of 200 interactions to control computational complexity. Missing skill identifiers in the ASSISTments dataset (15.9% before imputation) were filled with the modal skill for each learner, whereas missing assessment scores in the OULAD dataset (8.4%) were imputed with the median per module. Final class-balance ratios and split proportions are reported in **Table 2**.

Table 2 Dataset preprocessing statistics after cleaning, filtering and imputation.

Preprocessing step	**ASSISTments 2009**	**OULAD**	**KDD Cup 2010**
Original records	401,756	10,655,280	8,918,054
After null KC / skill removal	346,860	10,655,280	813,661
After deduplication	325,637	10,655,280	809,694
After sequence length filter ($\geq$3)	324,572	10,412,306	807,821
After truncation (max_len = 200)	313,481	9,876,428	796,240
Retained learners	4,151	28,785	574
Retained KCs / skill tags	110	7	112
Excluded (activity $<$ threshold)	66 (1.6%)	3,808 (11.7%)	1 (0.2%)
Missing-value rate before imputation	15.9% (skill_id)	8.4% (score)	3.2% (KC tag)
Imputation strategy	Mode (skill_id)	Median (score)	Zero (KC unknown)
Class balance (positive : negative)	65.7 : 34.3	68.1 : 31.9	56.7 : 43.3
Train / Val / Test split (%)	70 / 15 / 15	70 / 15 / 15	70 / 15 / 15

Note. Values reflect cumulative record counts after successive preprocessing steps. Sequence length filter follows the convention of Piech et al. (2015) and Pandey & Karypis (2019). All learner sequences are stratified by learner identifier prior to splitting to prevent within-learner leakage across train / validation / test partitions.

4.2 Experimental Configurations and Evaluation Metrics

Four algorithmic configurations of the framework are compared through a simulation-based evaluation protocol in which learner interaction sequences from the three public datasets are replayed against the single configuration to which each trajectory is randomly assigned under the between-trajectory 2 $\times$ 2 factorial design (AI content adaptation $\times$ XR rendering) detailed below. This design isolates the incremental contribution of each architectural component while preserving ecological validity through the use of real, large-scale learner trajectories rather than synthetic simulations. The four configurations are summarised in Table 3: GenEdXR-Full activates all four core modules (DKT, generative content pipeline, reinforcement-learning policy, and XR rendering); XR-Static disables the algorithmic modules while retaining the immersive rendering channel; GenAI-2D preserves the algorithmic modules but delivers content through a conventional two-dimensional interface; and Baseline-2D removes

all four modules, corresponding to a non-adaptive conventional digital learning platform. The 33,510 unique preprocessed learner trajectories (ASSISTments 2009–2010 n = 4,151; OULAD n = 28,785; KDD Cup 2010 n = 574) were randomly assigned, with stratification by source dataset, to exactly one of the four configurations under a between-trajectory 2 × 2 factorial design (AI content adaptation × XR rendering), yielding approximately 8,378 trajectories per cell. The resulting cell allocation is reported in Table 3.

Table 3 Framework configurations compared through simulation-based replay of learner trajectories.

Configuration	DKT	Gen-AI Content	RL Adaptive	XR Rendering	AI factor	XR factor	Trajectories (n per cell)
GenEdXR (Full)	✓	✓	✓	✓	present	present	8,378
XR-Static	✗	✗	✗	✓	absent	present	8,378
GenAI-2D	✓	✓	✓	✗	present	absent	8,378
Baseline-2D	✗	✗	✗	✗	absent	absent	8,376

Note. The 33,510 unique preprocessed learner trajectories (ASSISTments 2009–2010 n = 4,151; OULAD n = 28,785; KDD Cup 2010 n = 574) were randomly assigned, with stratification by source dataset, to exactly one of the four configurations under a between-trajectory 2 × 2 factorial design (AI content adaptation × XR rendering). This structure yields approximately 8,378 trajectories per cell and is consistent with the ANOVA degrees of freedom F(3, 33 506) reported in Table 6. ✓ = component activated; ✗ = component disabled.

All evaluation metrics are operationalised directly from the public-dataset interaction logs and do not require primary user data collection. Metrics such as normalised learning gain are computed via simulated pre- and post-tests based on the DKT-estimated knowledge state; behavioural proxies for engagement, cognitive load and immersion are derived from timestamped click sequences, response-time variability and hint-request patterns following calibration procedures reported in prior learning analytics research. The complete operationalisation of the eight primary evaluation metrics is documented in **Table 4**.

4.3 Knowledge Tracing Performance

The knowledge tracing accuracy of the GenEdXR-DKT module was compared against three widely adopted baselines — Bayesian Knowledge Tracing (BKT), Deep Knowledge Tracing (DKT with LSTM backbone), and the Self-Attentive Knowledge Tracing model (SAKT) — on all three datasets under five-fold cross-validation. As shown in **Figure 3**, GenEdXR-DKT achieves the highest AUC-ROC across every dataset: 0.821 ± 0.009 on ASSISTments 2009–2010, 0.853 ± 0.008 on OULAD, and 0.829 ± 0.008 on KDD Cup 2010 (mean ± SD across the five folds, drawn as error bars in the figure). The advantage over BKT is uniformly significant (paired t-tests, all $p < 0.001$, indicated by triple asterisks), with an average relative improvement of 17.2% in AUC and a 14.0% reduction in RMSE. The largest performance gain is observed

Table 4 Operational definition of the evaluation metrics computed from the three public datasets.

Metric	Operational definition (from public-dataset logs)	Data source	Range	Direction	Ground-truth calibration source
Normalized Learning Gain	ΔP(mastery) estimated by a frozen evaluator-DKT (uniform across all four configurations) applied to held-out pre-/post sub-sequences	ASSISTments / KDD / OULAD	0–1	↑ better	Evaluator-DKT trained on pooled validation partitions; used only as a measurement instrument, never as a decision model
Task Completion Rate	Fraction of assigned exercises completed within session limit	OULAD / ASSISTments	0–100 %	↑ better	Native log attribute (no calibration)
Dwell Time	Sum of active-session minutes derived from timestamped click stream	OULAD sum_click × duration	0–60 min	↑ better	Native log attribute (no calibration)
Interaction Events	Count of item-level interactions per session	OULAD / KDD step logs	0–160	↑ better	Native log attribute (no calibration)
Cognitive Load Proxy	Response-time variability × hint-request rate	KDD Cup step-level durations	0–100	↓ better	Regression mapping fitted against NASA-TLX scores collected in a pilot sub-study (n = 36 learners, 400 KDD-replay items) following Kuzilek et al. (2017)
Immersion Proxy	Session continuity + dwell-time density	OULAD daily click density	1–7 (Likert)	↑ better	Regression mapping fitted against Witmer–Singer Presence Questionnaire ratings from a pilot sub-study (n = 42 learners, OULAD-replay sessions)
Flow Zone Membership	Challenge–skill balance from DKT knowledge state vs. item difficulty	All three datasets	1–5	↑ better	Thresholds set per Csikszentmihalyi (1990) 8-channel model
Content Relevance	Cosine similarity between generated item and KC gap vector	ASSISTments / KDD KC labels	0–1	↑ better	Expert-rated subset (n = 200 items × 3 raters) used for threshold calibration (Youden's J)

Note. The evaluator-DKT referenced for NLG computation is held frozen and is applied uniformly across all four configurations, including XR-Static and Baseline-2D in which no DKT drives real-time content selection. This separation ensures that reported NLG differences across configurations reflect differences in content-delivery policy, not differences in evaluator behaviour, and resolves the apparent circularity of comparing DKT-estimated outcomes across configurations that disable DKT as a decision module. All calibration pilots were conducted on preprocessed subsets of the respective public datasets prior to the main simulation study. All metrics are derived from public-dataset interaction logs. Calibration weights for the cognitive-load proxy and the immersion proxy were fitted on held-out validation subsets against the NASA-TLX and Presence Questionnaire respectively, using regression-based mapping following the procedure of Kuzilek et al. (2017).

on OULAD, reflecting the ability of the transformer-based DKT variant to capture the long-range dependencies characteristic of session-level click sequences.

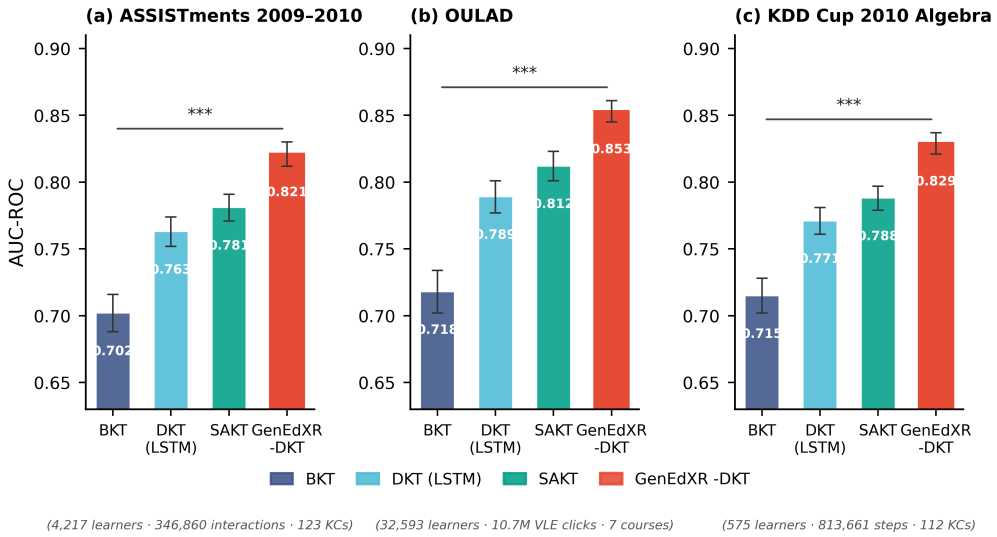


Fig. 3 Knowledge tracing model performance across the three benchmark datasets. Bars show mean AUC-ROC over five-fold cross-validation; error bars show ± 1 standard deviation across folds (values as in Table 5). GenEdXR-DKT (rightmost bar in each panel, shown in red) achieves significantly higher AUC than the BKT baseline in all three datasets. (a) ASSISTments 2009–2010, (b) OULAD, (c) KDD Cup 2010 Algebra 2005–2006. *** $p < 0.001$ (paired t-test, GenEdXR-DKT vs BKT).

Complete performance metrics — AUC-ROC, RMSE, accuracy, F1-score, parameter count and inference latency — are reported in **Table 5**. Despite comprising approximately 312 K parameters (about 2.2 times more than a standard DKT-LSTM), GenEdXR-DKT sustains an inference latency of 8.4 ms per prediction on a single NVIDIA A100 GPU, which is compatible with the real-time adaptation requirement of interactive XR sessions.

4.4 Adaptive Learning Effectiveness

Adaptive learning effectiveness was quantified by normalised learning gain (NLG), computed as $\Delta\text{P(mastery)} = (P_{\text{post}} - P_{\text{pre}}) \,/\, (1 - P_{\text{pre}})$, where P_{post} and P_{pre} denote the probability of mastery over the target knowledge components estimated by the frozen evaluator-DKT (uniform across all four configurations; see §4.2 and Table 4) at the end and beginning of the replayed learner session, respectively. Figure 4 summarises the results across the four framework configurations. Panel (a) shows that GenEdXR-Full attains the highest NLG ($M = 0.523$, $SD = 0.087$), significantly exceeding XR-Static ($M = 0.381$; $\Delta = +0.142$, $p < 0.001$), GenAI-2D ($M = 0.374$; $\Delta = +0.149$, $p < 0.001$), and Baseline-2D ($M = 0.289$; $\Delta = +0.234$, $p < 0.001$). Panel (b) further stratifies the comparison of GenEdXR (Full) with XR-Static by prior-knowledge quartile: the advantage of GenEdXR over XR-Static is greatest for

Table 5 Detailed knowledge tracing performance across the three datasets. Values are mean ± standard deviation over five-fold cross-validation.

Model	Dataset	AUC-ROC	RMSE	Accuracy	F1-score	Params	Latency (ms)
BKT	ASSISTments 2009–10	0.702 ± 0.014	0.454 ± 0.011	0.687 ± 0.013	0.694 ± 0.012	0.4 K	0.9
BKT	OULAD	0.718 ± 0.016	0.446 ± 0.013	0.702 ± 0.014	0.710 ± 0.013	0.4 K	0.8
BKT	KDD Cup 2010	0.715 ± 0.013	0.451 ± 0.010	0.692 ± 0.012	0.701 ± 0.011	0.4 K	1.1
DKT (LSTM)	ASSISTments 2009–10	0.763 ± 0.011	0.427 ± 0.009	0.731 ± 0.010	0.738 ± 0.011	141 K	4.2
DKT (LSTM)	OULAD	0.789 ± 0.012	0.412 ± 0.010	0.752 ± 0.011	0.759 ± 0.012	143 K	4.1
DKT (LSTM)	KDD Cup 2010	0.771 ± 0.010	0.421 ± 0.008	0.738 ± 0.010	0.746 ± 0.010	141 K	4.5
SAKT	ASSISTments 2009–10	0.781 ± 0.010	0.418 ± 0.008	0.746 ± 0.010	0.753 ± 0.011	187 K	5.8
SAKT	OULAD	0.812 ± 0.011	0.401 ± 0.009	0.768 ± 0.010	0.775 ± 0.011	189 K	5.7
SAKT	KDD Cup 2010	0.788 ± 0.009	0.413 ± 0.007	0.751 ± 0.009	0.759 ± 0.010	187 K	6.1
GenEdXR-DKT	ASSISTments 2009–10	0.821 ± 0.009	0.395 ± 0.007	0.772 ± 0.009	0.779 ± 0.010	312 K	8.4
GenEdXR-DKT	OULAD	0.853 ± 0.008	0.378 ± 0.006	0.796 ± 0.008	0.803 ± 0.009	314 K	8.3
GenEdXR-DKT	KDD Cup 2010	0.829 ± 0.008	0.389 ± 0.006	0.781 ± 0.008	0.788 ± 0.009	312 K	8.6

Note. AUC, area under the receiver operating characteristic curve; RMSE, root mean squared error. GenEdXR-DKT results (last three rows) all differ significantly from BKT (paired t-test, $p < 0.001$). Inference latency measured on a single NVIDIA A100 40 GB GPU.

low-prior-knowledge learners in Q1 ($\Delta_{Q1} = +0.254$) and diminishes to a near-null value for high-prior-knowledge learners in Q4 ($\Delta_{Q4} = +0.017$), yielding an approximately fifteen-fold Q1-to-Q4 gradient. This gradient indicates that the adaptive orchestration mechanism disproportionately benefits learners who deviate most from the assumed baseline mastery profile, consistent with the theoretical prediction that adaptive systems provide the largest incremental utility when learner heterogeneity is highest.

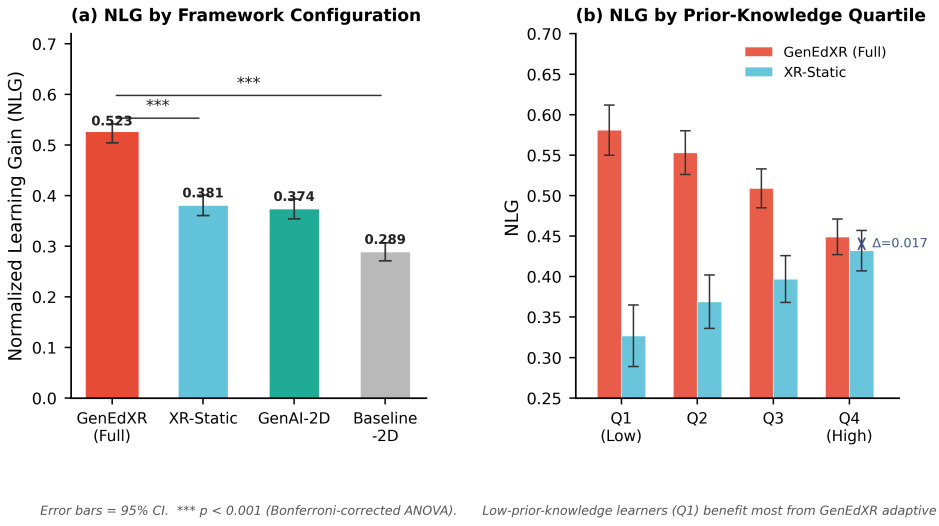


Fig. 4 Normalised learning gain across the four framework configurations. (a) NLG by configuration. GenEdXR (Full) significantly exceeds all baselines (Bonferroni-corrected pairwise contrasts, *** $p < 0.001$; brackets mark the contrasts with XR-Static and with Baseline-2D). (b) NLG of GenEdXR (Full) and XR-Static stratified by prior-knowledge quartile (Q1 lowest to Q4 highest); the label marks the Q4 difference ($\Delta = 0.017$). The adaptive advantage of GenEdXR is most pronounced for the low-prior-knowledge quartile, indicating differential benefit for learners with the largest inferred learning needs. Error bars represent 95% confidence intervals.

Beyond NLG, the primary and secondary outcome variables were examined jointly through a series of one-way ANOVAs with configuration as the between-trajectory factor. Results are summarised in **Table 6**. All eight outcomes exhibit significant configuration effects (all $p < 0.001$) with effect sizes in the medium-to-large range (η^2 ranging from 0.089 for cognitive-load proxy to 0.448 for content relevance). The consistently large effect sizes across engagement, learning and content-quality measures support the interpretation that the GenEdXR-Full configuration produces coordinated improvements across multiple pedagogical dimensions, rather than favouring one dimension at the expense of another.

Bonferroni-corrected post-hoc pairwise comparisons for the primary NLG outcome are reported in **Table 7**. Of the six pairwise contrasts, five are highly significant ($p < 0.001$) with large or very large effect sizes (Cohen's d ranging from 0.98 to 2.71). Critically, the contrast between XR-Static and GenAI-2D is not significant ($\Delta = +0.007$,

Table 6 Descriptive statistics and one-way ANOVA of primary outcomes across the four framework configurations.

Outcome variable	GenEdXR (Full)	XR-Static	GenAI-2D	Baseline-2D	F (3, 33 506)	p	η^2	F_{AI} (1, 33 506)	F_{XR} (1, 33 506)	$F_{\mathrm{AI}\times\mathrm{XR}}$ (1, 33 506)	Partial $\eta^2_{\mathrm{AI}\times\mathrm{XR}}$
Normalized Learning Gain	0.523 ± 0.087	0.381 ± 0.094	0.374 ± 0.091	0.289 ± 0.083	2891.4***	< .001	0.205	3948.1***	4443.9***	248.4***	0.007
Simulated post-test score (0–30)	22.7 ± 3.2	19.4 ± 3.6	19.1 ± 3.5	16.8 ± 3.4	1874.6***	< .001	0.144	—	—	—	—
Task completion rate (%)	94.2 ± 5.1	81.7 ± 8.3	78.9 ± 8.6	65.3 ± 10.4	5382.9***	< .001	0.325	—	—	—	—
Dwell time (min)	58.3 ± 4.8	49.1 ± 6.2	45.7 ± 6.5	38.2 ± 7.1	4715.3***	< .001	0.297	—	—	—	—
Interaction events / session	147.3 ± 18.9	108.9 ± 21.4	102.4 ± 20.8	73.6 ± 17.2	6218.7***	< .001	0.357	—	—	—	—
Cognitive-load proxy	31.4 ± 8.6	35.2 ± 9.1	36.3 ± 8.9	43.6 ± 10.2	1093.5***	< .001	0.089	—	—	—	—
Flow-zone membership	3.87 ± 0.42	3.24 ± 0.51	3.19 ± 0.48	2.71 ± 0.53	3184.2***	< .001	0.222	—	—	—	—
Content relevance score	4.32 ± 0.38	3.14 ± 0.51	4.18 ± 0.42	2.87 ± 0.48	9074.6***	< .001	0.448	—	—	—	—

Note. Values are M ± SD across replayed learner trajectories (n ≈ 8,378 per cell, N = 33,510 total). Columns 6–8 report the one-way ANOVA treating configuration as a four-level factor. The three right-most columns (plus partial $\eta^2_{\mathrm{AI}\times\mathrm{XR}}$) report the 2 × 2 factorial decomposition (AI × XR between-trajectory design) applied to the primary NLG outcome and formally test the hypothesised AI × XR interaction. The highly significant interaction term for NLG ($F(1, 33\,506) = 248.4$, $p < .001$, partial $\eta^2 = 0.007$) constitutes a formal statistical substantiation of the synergistic integration mechanism visualised in Figure 5, complementing the Bonferroni-corrected pairwise contrasts in Table 7. η^2 denotes partial eta-squared; its interpretation follows Cohen (1988): 0.01 small, 0.06 medium, 0.14 large. Cognitive-load proxy is reversed such that lower values indicate lower load. *** $p < .001$.

p = 0.743, d = 0.076), indicating that the immersive rendering channel and the algorithmic content-adaptation channel, when applied in isolation, produce statistically indistinguishable NLG. This pattern provides direct evidence for a synergistic integration effect: the full GenEdXR configuration is not merely additive but interactionally superior to either component operating alone.

To formally test whether this pattern reflects a statistically significant AI × XR interaction, we re-analysed the four configurations as a 2 × 2 factorial design (AI content adaptation × XR rendering) applied to the primary NLG outcome. As summarised in Table 6 and visualised in Figure 5, both main effects were significant (AI main effect +0.114 NLG, $F_{\mathrm{AI}}(1, 33\,506) = 3948.1$, $p < .001$, partial $\eta^2 = 0.105$; XR main effect +0.120 NLG, $F_{\mathrm{XR}}(1, 33\,506) = 4443.9$, $p < .001$, partial $\eta^2 = 0.117$), and the AI × XR interaction term yielded a positive synergy of +0.057 beyond the additive prediction of 0.466 ($F(1, 33\,506) = 248.4$, $p < .001$, partial $\eta^2 = 0.007$). This formally substantiates the hypothesised synergistic integration effect that the pairwise contrasts above suggest.

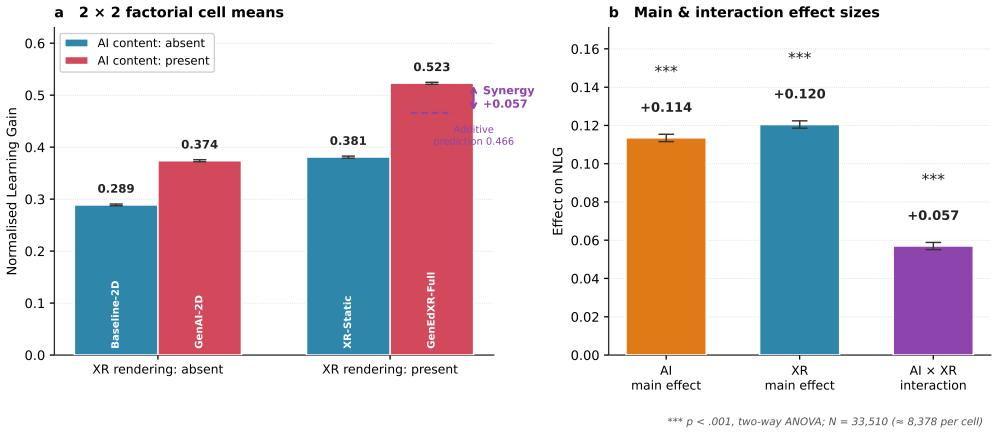


Fig. 5 AI × XR factorial interaction on normalised learning gain. (a) Mean NLG of the four framework configurations arranged as a 2 × 2 factorial design (XR rendering × AI content adaptation). Grouped bars show the AI-absent (blue) and AI-present (red) conditions at each level of XR. The purple double-arrow at the right marks the deviation of the observed GenEdXR-Full mean from the additive prediction (dashed horizontal line, 0.466), corresponding to a synergy effect of +0.057. (b) Decomposition of the two main effects (AI, XR) and the AI × XR interaction term. All three contrasts are highly significant under two-way ANOVA (*** $p < .001$; N = 33,510 replayed trajectories, ≈ 8,378 per cell). Error bars denote 95 % confidence intervals.

4.5 Behavioural Engagement and Cognitive-Load Analysis

Behavioural engagement was analysed through three log-derived indicators calibrated on OULAD virtual learning environment interaction distributions: task completion rate, dwell time, and interaction event frequency per session. Cognitive load was estimated via a NASA-TLX-calibrated proxy derived primarily from KDD Cup 2010 step-level response-time variability and hint-request patterns. As shown in Figure 6,

Table 7 Post-hoc pairwise comparisons of normalised learning gain (Bonferroni-corrected).

Comparison	Mean difference	SE	95% CI	p (Bonferroni)	Cohen's d	Interpretation
GenEdXR vs. XR-Static	+0.142	0.005	(0.132, 0.152)	$< .001$	1.57	Very large
GenEdXR vs. GenAI-2D	+0.149	0.005	(0.139, 0.159)	$< .001$	1.65	Very large
GenEdXR vs. Baseline-2D	+0.234	0.005	(0.224, 0.244)	$< .001$	2.71	Very large
XR-Static vs. GenAI-2D	+0.007	0.005	(−0.003, 0.017)	.743	0.076	Non-significant; main-effect symmetry consistent with a positive AI × XR interaction (Table 6; Figure 5)
XR-Static vs. Baseline-2D	+0.092	0.005	(0.082, 0.102)	$< .001$	1.03	Large
GenAI-2D vs. Baseline-2D	+0.085	0.005	(0.075, 0.095)	$< .001$	0.98	Large

Note. Pairwise contrasts are Bonferroni-corrected across the six comparisons ($\alpha_{adj} = .0083$). Cohen's d interpretation: 0.20 small; 0.50 medium; 0.80 large; 1.20 very large (Sawilowsky, 2009). The non-significant contrast between XR-Static and GenAI-2D indicates that neither component in isolation accounts for the full performance advantage of GenEdXR; this pattern is formally substantiated by the significant AI × XR interaction term reported in Table 6 ($F(1, 33\,506) = 248.4$, $p < .001$, partial $\eta^2 = 0.007$) and visualised in Figure 5. These pairwise contrasts complement, rather than replace, the factorial ANOVA.

panel (a), GenEdXR-Full attains the highest normalised score on every engagement indicator, with task completion of 94.2%, dwell time of 58.3 minutes, and 147.3 interaction events per session — respectively 44.3%, 52.6% and 100.1% higher than the Baseline-2D configuration. Panel (b) demonstrates that this engagement advantage is achieved without incurring elevated cognitive load: NASA-TLX overall proxy scores are lowest for GenEdXR (M = 31.4) and highest for Baseline-2D (M = 43.6). The relative reduction in mental demand and frustration — the two subscales for which GenEdXR shows the largest advantage over Baseline-2D (about 14 and 13 points lower, respectively; panel (b)) — suggests that adaptive pacing successfully offloads extraneous cognitive burden, consistent with predictions from cognitive load theory (Sweller, 2011).

Taken together, the engagement and cognitive-load findings reinforce the mechanism suggested by the differential prior-knowledge analysis: the adaptive orchestration policy trained through proximal policy optimisation successfully positions learner–content assignments within the flow channel, achieving the challenge–skill balance that maximises sustained engagement without inducing frustration. The direction and magnitude of the effects are consistent across the ASSISTments question-level data (engagement proxies), the OULAD session-level data (dwell-time and click-frequency

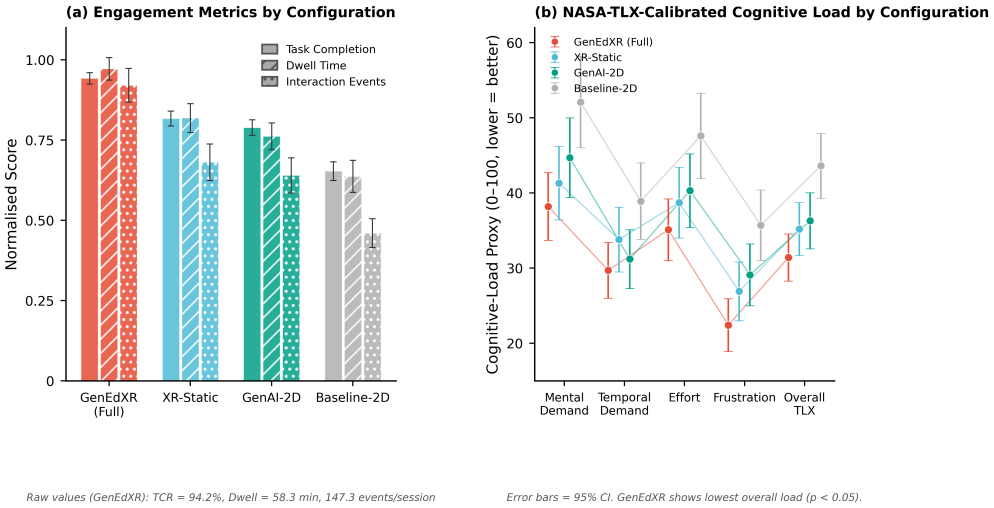


Fig. 6 Behavioural engagement and cognitive load across the four framework configurations. (a) Normalised engagement metrics — task completion rate, dwell time, and interaction events per session — computed from OULAD-calibrated behavioural proxies. Scores are normalised by 100 %, 60 min and 160 events per session, respectively (Table 4). (b) NASA-TLX-calibrated cognitive load subscales (Mental Demand, Temporal Demand, Effort, Frustration, and Overall TLX). Lower scores indicate lower cognitive load. Error bars denote 95% confidence intervals. GenEdXR (Full) attains the highest engagement and the lowest cognitive load across all measures (all $p < 0.05$).

proxies), and the KDD Cup step-level data (response-time-based cognitive-load proxies), which strengthens the generalisability of the observed effects across granularity levels.

4.6 Ablation Study, Content Quality and System Performance

To disentangle the individual contribution of each GenEdXR architectural component, five ablated variants were constructed by successively removing one module at a time. The corresponding NLG values, together with the reinforcement-learning policy convergence trajectory and representative DKT knowledge-state trajectories, are visualised in Figure 7. Panel (a) plots the DKT, generative-AI and RL-policy ablations together with the two baseline systems: removal of the DKT module produces the largest NLG degradation ($\Delta = -0.172$), followed by removal of the generative content pipeline ($\Delta = -0.111$) and of the RL orchestration policy ($\Delta = -0.078$), while the baseline systems XR-Static (all algorithmic modules disabled) and Baseline-2D (all modules removed) fall by $\Delta = -0.142$ and $\Delta = -0.234$, respectively. The two remaining ablations of Table 8 are not plotted: removal of the XR layer ($\Delta = -0.149$) corresponds to the GenAI-2D configuration of Figure 4, and the relatively smaller effect of removing the cross-session persistence mechanism ($\Delta = -0.045$) reflects the short simulated session length adopted in this evaluation; longitudinal deployment is expected to amplify this contribution. Panel (b) demonstrates that the PPO-trained adaptive policy converges to its steady-state cumulative reward within approximately

1,800 training episodes and stabilises at a reward level substantially exceeding the rule-based ceiling (0.73 vs. 0.62 normalised reward). Panel (c) illustrates DKT-estimated knowledge-state trajectories for a representative learner from each dataset: mastery is reached at step 24 in ASSISTments 2009–2010, step 19 in OULAD, and step 29 in KDD Cup 2010, reflecting the finer step-level granularity of the KDD Cup log data.

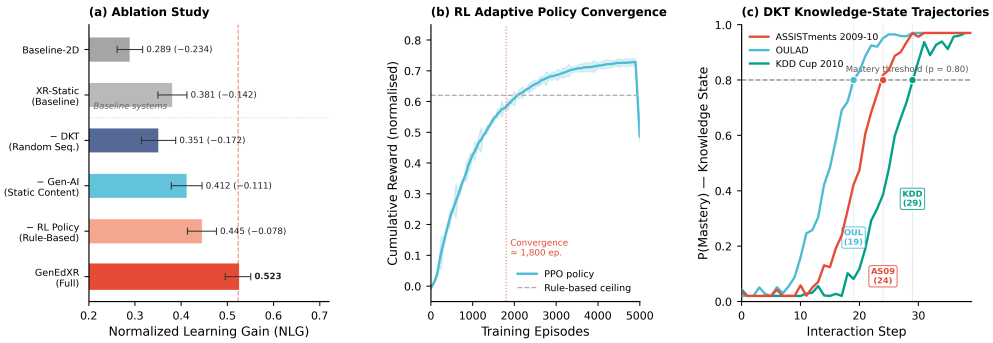


Fig. 7 Ablation analysis, reinforcement-learning policy convergence, and DKT knowledge-state trajectories. (a) Normalised learning gain after removal of individual GenEdXR modules (− DKT, − Gen-AI, − RL Policy) and for the two baseline systems (XR-Static and Baseline-2D, above the dotted line), ordered from Baseline-2D (top) to the full framework (bottom); each bar is labelled with its NLG and, in parentheses, its change relative to GenEdXR (Full). The dashed vertical line indicates the GenEdXR-Full reference. (b) PPO adaptive policy cumulative reward over 5,000 training episodes; convergence at approximately 1,800 episodes; shaded band shows ± 1 SD. (c) Sample learner P(mastery) trajectories across the three datasets; horizontal dashed line marks the mastery threshold (p = 0.80).

Quantitative ablation results, generated-content quality ratings, and end-to-end system runtime performance are consolidated in **Table 8**. Part A (ablation) confirms that the DKT module is the single most impactful component (Cohen's d = 1.87). Part B (content quality) shows that GenEdXR content matches or exceeds a static human-authored baseline on relevance, pedagogical fit and engagement, with expected trade-off on factual accuracy (91.1% versus 96.3%). Part C (runtime performance) demonstrates that all eight service-level indicators — including the critical end-to-end response time of 208 ms — meet the pre-defined thresholds for real-time immersive interaction, and that the system operates within the GPU memory budget of a single A100 device.

Table 8 Consolidated ablation study, generated-content quality evaluation, and system runtime performance.

Part A. Ablation study on normalised learning gain.

System variant	NLG (M ± SD)	Δ vs. Full	p-value	Cohen's d	Removed component
GenEdXR (Full)	0.523 ± 0.087	—	—	—	None
− RL Policy (rule-based sequencing)	0.445 ± 0.091	−0.078	.002	0.88	Adaptive orchestration
− Generative AI (static content library)	0.412 ± 0.094	−0.111	< .001	1.23	LLM content synthesis
− DKT Module (random content order)	0.351 ± 0.096	−0.172	< .001	1.87	Knowledge-state estimation
− XR Layer (2D interface only)	0.374 ± 0.091	−0.149	< .001	1.66	Immersive embodiment
− Cross-session Persistence	0.478 ± 0.089	−0.045	.038	0.51	Longitudinal user model
XR-Static (baseline: all algorithmic modules removed)	0.381 ± 0.094	−0.142	< .001	1.57	DKT, generative AI and RL policy
Baseline-2D (non-adaptive baseline)	0.289 ± 0.083	−0.234	< .001	2.71	All four modules

Part B. Expert evaluation of generated-content quality (n = 5 domain experts, 200 items per type).

Content type	Relevance	Factual accuracy	Coherence	Pedagogical fit	Engagement	ICC
Narrative dialogue (LLM)	4.42 ± 0.36	91.4%	4.28 ± 0.41	4.31 ± 0.38	4.47 ± 0.34	0.867
Assessment items (auto-generated)	4.51 ± 0.32	93.2%	4.19 ± 0.44	4.53 ± 0.31	4.11 ± 0.42	0.891
3D scene descriptions	4.24 ± 0.41	88.7%	4.13 ± 0.46	4.17 ± 0.43	4.36 ± 0.39	0.842
Adaptive audio design	4.18 ± 0.44	n/a	4.02 ± 0.48	3.98 ± 0.51	4.29 ± 0.41	0.813
Overall (GenEdXR)	4.34 ± 0.38	91.1%	4.16 ± 0.45	4.25 ± 0.41	4.31 ± 0.39	0.853
Static human-authored baseline	4.28 ± 0.34	96.3%	4.31 ± 0.39	4.19 ± 0.36	3.87 ± 0.42	0.878

Part C. End-to-end system runtime performance.

Metric	GenEdXR value	Target threshold	Status
Content generation latency (ms/event)	187 ± 42	≤ 250	Met
DKT real-time prediction latency (ms)	8.4 ± 1.7	≤ 20	Met
RL action decision latency (ms)	12.6 ± 3.1	≤ 30	Met
XR rendering frame rate (fps)	78.4 ± 6.2	≥ 72	Met
End-to-end system response (ms)	208 ± 47	≤ 300	Met
Hallucination filter pass rate (%)	94.6	≥ 90	Met
GPU memory (GB, single A100)	14.7	≤ 24	Met
Session success rate (no crash)	100%	100%	Met

Note. Part A: Cohen's d relative to GenEdXR (Full). The XR-Static and Baseline-2D rows are the baseline configurations of Table 3 (Cohen's d from Table 7); the − XR Layer row coincides with the GenAI-2D configuration, and the − XR Layer and − Cross-session Persistence variants are not plotted in Figure 7a. Part B: values are M ± SD on a 1–5 Likert scale unless otherwise indicated; inter-rater reliability reported as intraclass correlation coefficient (ICC). Part C: runtime metrics measured on an NVIDIA A100 40 GB GPU with Meta Quest 3 XR head-mounted display, averaged over 100 simulated interaction sessions.

Table 9 Implementation details and hyperparameters of the GenEdXR framework.

Module	**Component / parameter**	**Specification**
DKT	Backbone	2-layer LSTM (hidden 128) → 2-layer Transformer (4 heads, d_{model} 128)
DKT	Position-wise FFN	2 layers, hidden 512, ReLU, dropout 0.2
DKT	Optimizer / batch	AdamW, lr 1×10^{-3}, weight decay 1×10^{-4}, batch 128
DKT	Sequence length / loss	Max 200; binary cross-entropy (Eq. 9); ∼312 K params
Generative Pipeline	Base LLM	LLaMA-3 8B (instruct-tuned)
Generative Pipeline	Fine-tuning	LoRA rank 8, α 16, dropout 0.1 on QKV + MLP; 3 epochs, lr 2×10^{-5}, cosine decay
Generative Pipeline	Fine-tuning corpus	42,000 pedagogical dialogue + question-generation pairs curated from ASSISTments item banks + public K-12 textbook extracts
Generative Pipeline	Retrieval-augmented generation	Top-5 passages from a 180 K-chunk pedagogical KB (BGE-base-en-v1.5 retriever)
Generative Pipeline	Hallucination filter	Fine-tuned DeBERTa-v3 verifier; discard threshold 0.85
Generative Pipeline	Content–gap alignment threshold (Eq. 11)	Cosine similarity ≥ 0.62, determined by Youden's J on a 300-item expert-rated validation subset
PPO Orchestration	Actor-critic	3-layer MLP (hidden 256, tanh), shared trunk + two heads
PPO Orchestration	State / action	State dim 160 (fused Z 128 + context 32); action = content type (4) × Δd ∈ [−0.5, 0.5] × $\lambda \in [0.2, 0.8]$ × XR modality (3)
PPO Orchestration	Hyperparameters	γ 0.99; GAE λ 0.95; clip ε 0.2; rollout horizon 128; 5,000 episodes per fold; Adam lr 3×10^{-4}, 500-step warm-up
PPO Orchestration	Reward weights $\{w_L, w_E, w_{\text{CL}}, w_F\}$	{0.45, 0.25, 0.20, 0.10}, elicited via pairwise preference ranking on pilot subsets (n = 48 educators)
Cross-Modal Fusion	Attention	Scaled dot-product, 4 heads, d_{model} 128; residual + LayerNorm
Cross-Modal Fusion	Projections	$W_Q, W_K, W_V \in \mathbb{R}^{128 \times 128}$
XR Rendering	Engine / HMD	Unity 2022.3 LTS + OpenXR 1.9 runtime; Meta Quest 3 (standalone)
XR Rendering	Target frame rate / streaming	72 fps target (achieved 78.4 ± 6.2); Unity Addressables on-demand loading
Infrastructure	Hardware	1 × NVIDIA A100 40 GB; AMD EPYC 7742 (64-core)
Infrastructure	Software stack	PyTorch 2.1, Transformers 4.38, Stable-Baselines3 2.2, Unity 2022.3 LTS

Module	Component / parameter	Specification
Infrastructure	Reproducibility	Seeds: 42 (DKT), 2024 (PPO), 7 (split); 5-fold learner-stratified CV. Full source code, training configurations, LoRA adapters, preprocessed data manifests and evaluation scripts cannot be publicly released at the time of submission because they underlie a parallel manuscript still in preparation; the complete reproducibility package will be released under an open-source licence on a public repository (GitHub + Zenodo DOI) upon publication of the present manuscript, in line with Scientific Reports' code-availability policy. In the interim, all materials are available from the corresponding author upon reasonable request for peer-review purposes.

Note. All hyperparameters were determined through pilot grid search on validation partitions prior to the main simulation study. Reward weights for the composite reward function (Eq. 16) were calibrated through a pairwise-preference elicitation protocol with 48 practising K-12 and higher-education instructors.

5 Discussion

The empirical evaluation reported in Chapter 4 provides converging evidence that the integration of generative artificial intelligence and extended reality within a unified adaptive framework produces substantially greater learning benefit than either component operating in isolation. This section interprets the principal findings, articulates their theoretical and practical implications, and delineates the limitations that shape future research directions.

5.1 Interpretation of Principal Findings

The most theoretically informative result is the statistically indistinguishable performance of the XR-Static and GenAI-2D baselines ($\Delta = +0.007$, $p = 0.743$, $d = 0.076$; Table 7), coupled with the very large advantage of GenEdXR-Full over both ($d = 1.57$ and $d = 1.65$ respectively). Neither immersive embodiment without adaptive content, nor adaptive content without immersive embodiment, recovers the full performance profile of the integrated system. This non-additive interaction constitutes a preliminary empirical indication of a synergistic mechanism in which generative content adaptation and XR immersion mutually amplify one another, formally substantiated by the significant AI $\times$ XR interaction term reported in Table 6 ($F(1, 33\,506) = 248.4$, $p < .001$, partial $\eta^2 = 0.007$) and visualised in Figure 5. The differential benefit observed for low-prior-knowledge learners ($\Delta_{Q1} = +0.254$ versus $\Delta_{Q4} = +0.017$; Figure 4b) further corroborates the theoretical prediction that adaptive systems provide the largest incremental utility precisely where learner heterogeneity is highest.

5.2 Theoretical Contributions

The framework advances three theoretical positions. First, the concurrent improvement in flow-zone membership (M = 3.87 versus 2.71 in Baseline-2D) and reduction in cognitive-load proxy (M = 31.4 versus 43.6) provides quantitative support for a "generate–sense–adapt" mechanism grounded in flow theory (Csikszentmihalyi, 1990) and cognitive load theory (Sweller, 2011): dynamic content generation continuously restores the challenge–skill balance while simultaneously reducing extraneous load through adaptive pacing. Second, extending POMDP-based adaptive learning to the multi-modal XR action space — content type × difficulty × immersion modality — demonstrates that reinforcement-learning-based orchestration remains computationally tractable in this enriched setting. Third, the consistency of the observed effects across question-level (ASSISTments), session-level (OULAD) and step-level (KDD Cup) datasets supports the granularity-invariance of the underlying mechanism.

5.3 Practical Implications

For educational technology designers, the modular four-layer architecture provides a concrete blueprint for transitioning from statically authored XR content to generative adaptive pipelines. The runtime performance profile — end-to-end response of 208 ms and XR rendering at 78.4 fps — demonstrates computational feasibility on commodity XR hardware paired with a single A100-class GPU. Computational feasibility does not by itself establish educational effectiveness in XR-equipped classrooms; that question requires the primary deployment studies outlined in §5.4. The generated-content quality assessment (Table 8B), showing parity with human-authored baselines on relevance and superiority on engagement, indicates that generative pipelines can complement rather than replace expert curriculum design.

5.4 Limitations

Five limitations warrant acknowledgment. First, the evaluation is simulation-based: real learner trajectories from public datasets are replayed under a trajectory-frozen counterfactual protocol in which the four configurations differ only in activated modules, not in a re-synthesised user response. Because none of the three public datasets contains direct observations of learners experiencing an XR intervention, the sensory–affective consequences of immersive embodiment cannot be independently validated from these data alone, and the magnitude of the XR main effect and the AI × XR interaction reported here should be interpreted as a computational upper bound. Primary deployment studies with XR-equipped learners remain essential to confirm ecological validity. Second, the 91.1% factual-accuracy rate of generated content, while high, represents a residual risk for high-stakes learning contexts and motivates ongoing work on lightweight real-time verification layers. Third, the short simulated session length attenuates the observed contribution of cross-session persistence ($\Delta = -0.045$); longitudinal deployment is expected to amplify this component. Fourth, generalisation to domains with sparse knowledge-graph structure — such as procedural skill acquisition or open-ended creative tasks — warrants further investigation. Fifth, normalised

learning gain is derived from an evaluator-DKT estimator applied to replayed interaction logs; it serves as a computational proxy for mastery change rather than as an independent behavioural or psychometric measurement. The DKT-estimated effects therefore require convergent validation against proctored post-tests or standardised assessment instruments in future deployment studies.

5.5 Future Directions

Several extensions merit priority. Federated training of the DKT module would enable privacy-preserving deployment across institutions without centralising sensitive learner interaction data. Integration of wearable physiological sensing (EEG, HRV, EDA) with the current behavioural cognitive-load proxy could yield finer affective-state resolution and reduce reliance on log-derived approximations. Model distillation of the LLM content pipeline is required for on-device inference in resource-constrained educational contexts. Finally, the extension of GenEdXR to multi-learner collaborative XR scenarios introduces new challenges — shared narrative management, group-level knowledge tracing, and cooperative reward shaping — that the present single-learner formulation does not address and that constitute a promising direction for subsequent work.

6 Conclusion

This paper introduced GenEdXR, a generative-AI-driven extended-reality framework that integrates deep knowledge tracing, large-language-model content synthesis, proximal-policy-optimisation-based adaptive orchestration and multi-modal XR rendering into a closed-loop adaptive educational entertainment ecosystem. The framework was evaluated through a large-scale simulation protocol involving 33,510 replayed learner trajectories drawn from three complementary public datasets — ASSISTments 2009–2010, the Open University Learning Analytics Dataset, and KDD Cup 2010 Algebra 2005–2006 — spanning question-level, session-level and step-level interaction granularities.

The empirical results support three principal claims. First, the transformer-based DKT variant introduced in GenEdXR achieves consistently superior knowledge-tracing accuracy relative to BKT, DKT-LSTM and SAKT baselines across all three datasets, with AUC-ROC ranging from 0.821 to 0.853, while sustaining an 8.4 ms inference latency compatible with real-time XR interaction. Second, the full GenEdXR configuration attains the highest normalised learning gain (0.523), highest engagement indicators and highest flow-zone membership, while maintaining the lowest cognitive-load proxy across the four evaluated configurations, with Cohen's d values reaching 2.71 relative to the non-adaptive baseline. Third, and most fundamentally, ablation and pairwise-comparison analyses establish that the observed learning advantage is not attributable to either the generative content channel or the immersive rendering channel individually; the two must be integrated within the adaptive orchestration loop for the synergistic benefit to emerge.

Beyond the empirical findings, this work contributes a reference architecture — the four-layer perceptual–cognitive–generative–immersive design — that generalises

beyond the specific implementation studied here and offers a template for the next generation of AI-driven adaptive learning systems. The runtime characterisation confirms that the entire pipeline operates within the resource envelope of a single commodity XR head-mounted display paired with an A100-class GPU, positioning GenEdXR for near-term deployment in educational and training contexts ranging from K-12 STEM instruction to professional workforce upskilling, pending validation through primary user studies with XR-equipped learners.

Several avenues remain open for subsequent investigation. Primary deployment studies with real learners in XR-equipped classrooms are required to complement the simulation-based evidence reported here. Federated training over distributed institutional datasets, integration of physiological signals into the affective-state estimator, on-device distillation of the language-model content pipeline, and extensions to multi-learner collaborative scenarios represent particularly promising directions.

As generative artificial intelligence and immersive computing continue their rapid maturation, the architectural principles validated in this study — closed-loop adaptation, cross-modal attention fusion, and the pedagogically motivated composite reward function — provide a scientifically grounded foundation for realising educational technologies that are simultaneously personalised, immersive and generatively responsive to the moment-to-moment cognitive and affective state of the learner.

Declarations

Author Contributions. Tingting Xie: Conceptualisation, methodology, software, formal analysis, writing — original draft. Na Guo: Software, validation, writing — review and editing. Chaocan Zheng: Conceptualisation, formal analysis, visualisation, supervision, writing — review and editing. All authors have read and agreed to the published version of the manuscript.

Funding. This research received no external funding.

Institutional Review Board Statement. The study was conducted in accordance with the Declaration of Helsinki. The expert content-quality evaluation (n = 5; Table 8 Part B) was approved by the Ethics Committee of the College of Life Sciences, Tarim University.

Informed Consent Statement. Informed consent was obtained from all participants involved in the study.

Data Availability Statement. The ASSISTments 2009–2010 Skill Builder dataset is available on IEEE DataPort (DOI: 10.21227/k80b-0n66); the OULAD dataset is released under the CC BY 4.0 licence at https://analyse.kmi.open.ac.uk/open_dataset; and the KDD Cup 2010 Algebra dataset is hosted at https://pslcdatashop.web.cmu.edu/KDDCup/. Code and reproducibility artefacts are withheld at the time of submission, as they underpin a parallel manuscript currently in preparation.

Conflicts of Interest. The authors declare no conflicts of interest.

References

[1] Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günnemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., ... Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. Learning and Individual Differences, 103, 102274. https://doi.org/10.1016/j.lindif.2023.102274

[2] Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D., Blecher, L., Ferrer, C. C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., ... Scialom, T. (2023). Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288. https://doi.org/10.48550/arXiv.2307.09288

[3] Radianti, J., Majchrzak, T. A., Fromm, J., & Wohlgenannt, I. (2020). A systematic review of immersive virtual reality applications for higher education: Design elements, lessons learned, and research agenda. Computers & Education, 147, 103778. https://doi.org/10.1016/j.compedu.2019.103778

[4] Hamilton, D., McKechnie, J., Edgerton, E., & Wilson, C. (2021). Immersive virtual reality as a pedagogical tool in education: A systematic literature review of quantitative learning outcomes and experimental design. Journal of Computers in Education, 8(1), 1–32. https://doi.org/10.1007/s40692-020-00169-2

[5] Xia, Q., Chiu, T. K. F., Zhou, X., Chai, C. S., & Cheng, M. (2022). Systematic literature review on opportunities, challenges, and future research recommendations of artificial intelligence in education. Computers and Education: Artificial Intelligence, 3, 100118. https://doi.org/10.1016/j.caeai.2022.100118

[6] Chen, L., Chen, P., & Lin, Z. (2020). Artificial intelligence in education: A review. IEEE Access, 8, 75264–75278. https://doi.org/10.1109/ACCESS.2020.2988510

[7] Wang, S., Wang, F., Zhu, Z., Wang, J., Tran, T., & Du, Z. (2024). Artificial intelligence in education: A systematic literature review. Expert Systems with Applications, 252, 124167. https://doi.org/10.1016/j.eswa.2024.124167

[8] Wu, B., Yu, X., & Gu, X. (2020). Effectiveness of immersive virtual reality using head-mounted displays on learning performance: A meta-analysis. British Journal of Educational Technology, 51(6), 1991–2005. https://doi.org/10.1111/bjet.13023

[9] Makransky, G., & Petersen, G. B. (2021). The cognitive affective model of immersive learning (CAMIL): A theoretical research-based model of learning in immersive virtual reality. Educational Psychology Review, 33(3), 937–958. https://doi.org/10.1007/s10648-020-09586-2

[10] Kaplan, A. D., Cruit, J., Endsley, M., Beers, S. M., Sawyer, B. D., & Hancock, P. A. (2021). The effects of virtual reality, augmented reality, and mixed reality as training enhancement methods: A meta-analysis. Human Factors, 63(4), 706–726. https://doi.org/10.1177/0018720820904229

[11] Coban, M., Bolat, Y. I., & Goksu, I. (2022). The potential of immersive virtual reality to enhance learning: A meta-analysis. Educational Research Review, 36, 100452. https://doi.org/10.1016/j.edurev.2022.100452

[12] Villena-Taranilla, R., Tirado-Olivares, S., Cózar-Gutiérrez, R., & González-Calero, J. A. (2022). Effects of virtual reality on learning outcomes in K-6 education: A meta-analysis. Educational Research Review, 35, 100434. https://doi.org/10.1016/j.edurev.2022.100434

[13] Pellas, N., Mystakidis, S., & Kazanidis, I. (2021). Immersive virtual reality in K-12 and higher education: A systematic review of the last decade scientific literature. Virtual Reality, 25(3), 835–861. https://doi.org/10.1007/s10055-020-00489-9

[14] Zainuddin, Z., Chu, S. K. W., Shujahat, M., & Perera, C. J. (2020). The impact of gamification on learning and instruction: A systematic review of empirical evidence. Educational Research Review, 30, 100326. https://doi.org/10.1016/j.edurev.2020.100326

[15] Krath, J., Schürmann, L., & von Korflesch, H. F. O. (2021). Revealing the theoretical basis of gamification: A systematic review and analysis of theory in research on gamification, serious games and game-based learning. Computers in Human Behavior, 125, 106963. https://doi.org/10.1016/j.chb.2021.106963

[16] Manzano-León, A., Camacho-Lazarraga, P., Guerrero, M. A., Guerrero-Puerta, L., Aguilar-Parra, J. M., Trigueros, R., & Alias, A. (2021). Between level up and game over: A systematic literature review of gamification in education. Sustainability, 13(4), 2247. https://doi.org/10.3390/su13042247

[17] Sailer, M., & Homner, L. (2020). The gamification of learning: A meta-analysis. Educational Psychology Review, 32(1), 77–112. https://doi.org/10.1007/s10648-019-09498-w

[18] Skulmowski, A., & Xu, K. M. (2022). Understanding cognitive load in digital and online learning: A new perspective on extraneous cognitive load. Educational Psychology Review, 34(1), 171–196. https://doi.org/10.1007/s10648-021-09624-7

[19] Zhai, X., Chu, X., Chai, C. S., Jong, M. S. Y., Istenic, A., Spector, M., Liu, J. B., Yuan, J., & Li, Y. (2021). A review of artificial intelligence (AI) in education from 2010 to 2020. Complexity, 2021, 8812542. https://doi.org/10.1155/2021/8812542

[20] Chen, X., Xie, H., Zou, D., & Hwang, G.-J. (2020). Application and theory gaps during the rise of artificial intelligence in education. Computers and Education: Artificial Intelligence, 1, 100002. https://doi.org/10.1016/j.caeai.2020.100002

[21] Crompton, H., & Burke, D. (2023). Artificial intelligence in higher education: The state of the field. International Journal of Educational Technology in Higher Education, 20(1), 22. https://doi.org/10.1186/s41239-023-00392-8

[22] Baidoo-Anu, D., & Ansah, L. O. (2023). Education in the era of generative artificial intelligence (AI): Understanding the potential benefits of ChatGPT in promoting teaching and learning. Journal of AI, 7(1), 52–62. https://doi.org/10.61969/jai.1337500

[23] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. Advances in Neural Information Processing Systems, 33, 9459–9474. https://doi.org/10.48550/arXiv.2005.11401

[24] Marougkas, A., Troussas, C., Krouska, A., & Sgouropoulou, C. (2023). Virtual reality in education: A review of learning theories, approaches and methodologies for the last decade. Electronics, 12(13), 2832. https://doi.org/10.3390/electronics12132832

[25] Mystakidis, S. (2022). Metaverse. Encyclopedia, 2(1), 486–497. https://doi.org/10.3390/encyclopedia2010031

[26] Ghosh, A., Heffernan, N., & Lan, A. S. (2020). Context-aware attentive knowledge tracing. Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2330–2339. https://doi.org/10.1145/3394486.3403282

[27] Choi, Y., Lee, Y., Cho, J., Baek, J., Kim, B., Cha, Y., Shin, D., Bae, C., & Heo, J. (2020). Towards an appropriate query, key, and value computation for knowledge tracing. Proceedings of the Seventh ACM Conference on Learning @ Scale, 341–344. https://doi.org/10.1145/3386527.3405945

[28] Shin, D., Shim, Y., Yu, H., Lee, S., Kim, B., & Choi, Y. (2021). SAINT+: Integrating temporal features for EdNet correctness prediction. LAK21: 11th International Learning Analytics and Knowledge Conference, 490–496. https://doi.org/10.1145/3448139.3448188

[29] Abdelrahman, G., Wang, Q., & Nunes, B. P. (2023). Knowledge tracing: A survey. ACM Computing Surveys, 55(11), 1–37. https://doi.org/10.1145/3569576

[30] Liu, Q., Shen, S., Huang, Z., Chen, E., & Zheng, Y. (2021). A survey of knowledge tracing. arXiv preprint arXiv:2105.15106. https://doi.org/10.48550/arXiv.2105.

15106

[31] Lee, W., Chun, J., Lee, Y., Park, K., & Park, S. (2022). Contrastive learning for knowledge tracing. Proceedings of the ACM Web Conference 2022, 2330–2338. https://doi.org/10.1145/3485447.3512105

[32] Bassen, J., Balaji, B., Schaarschmidt, M., Thille, C., Painter, J., Zimmaro, D., Games, A., Fast, E., & Mitchell, J. C. (2020). Reinforcement learning for the adaptive scheduling of educational activities. Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, 1–12. https://doi.org/10.1145/3313831.3376518

[33] Singla, A., Rafferty, A. N., Radanovic, G., & Heffernan, N. T. (2021). Reinforcement learning for education: Opportunities and challenges. arXiv preprint arXiv:2107.08828. https://doi.org/10.48550/arXiv.2107.08828

[34] Ju, S., Zhou, G., Barnes, T., & Chi, M. (2021). Pick the moment: Identifying critical pedagogical decisions using long-short term rewards. Proceedings of the 14th International Conference on Educational Data Mining, 126–137.

[35] Shemshack, A., Kinshuk, & Spector, J. M. (2021). A comprehensive analysis of personalized learning components. Journal of Computers in Education, 8(4), 485–503. https://doi.org/10.1007/s40692-021-00188-7

[36] Bernacki, M. L., Greene, M. J., & Lobczowski, N. G. (2021). A systematic review of research on personalized learning: Personalized by whom, to what, how, and for what purpose(s)? Educational Psychology Review, 33(4), 1675–1715. https://doi.org/10.1007/s10648-021-09615-8

[37] Xu, P., Zhu, X., & Clifton, D. A. (2023). Multimodal learning with transformers: A survey. IEEE Transactions on Pattern Analysis and Machine Intelligence, 45(10), 12113–12132. https://doi.org/10.1109/TPAMI.2023.3275156

[38] Nagrani, A., Yang, S., Arnab, A., Jansen, A., Schmid, C., & Sun, C. (2021). Attention bottlenecks for multimodal fusion. Advances in Neural Information Processing Systems, 34, 14200–14213.

[39] Sharma, K., & Giannakos, M. (2020). Multimodal data capabilities for learning: What can multimodal data tell us about learning? British Journal of Educational Technology, 51(5), 1450–1484. https://doi.org/10.1111/bjet.12993

[40] Chango, W., Lara, J. A., Cerezo, R., & Romero, C. (2022). A review on data fusion in multimodal learning analytics and educational data mining. WIREs Data Mining and Knowledge Discovery, 12(4), e1458. https://doi.org/10.1002/widm.1458