Research Article RJ-43

Researches on Multidisciplinary Approaches 2026 (2) ISSN:2791-9099

A Hybrid Quantum-Cognitive Framework for Personalized Learning: Simulation-Based Validation and System Design

Özel Sebetci / Prof. Dr.
Aydın Adnan Menderes University, Aydın Vocational College, Department of Computer Technologies, Computer Programming Programme
osebetci@adu.edu.tr

Abstract

The convergence of quantum cognition, quantum reinforcement learning, and neuroadaptive educational technologies has opened new possibilities for modeling learner variability beyond the limits of conventional adaptive systems. This study introduces the Quantum-Cognitive Adaptive Learning (QCAL) framework as a hybrid architecture that combines quantum-inspired cognitive state modeling, Quantum Reinforcement Learning (QRL) optimized via the Quantum Natural Gradient (QNG), and biometric-informed adaptive feedback.

Within QCAL, each learner's cognitive state is represented as a probabilistic quantum-inspired wavefunction (ψ), allowing the system to model conceptual uncertainty, contextuality, and dynamic transitions between partially competing knowledge states. A meta-adaptive reward correction mechanism incorporates simulated EEG- and eye-tracking-derived signals to regulate task difficulty, pacing, and instructional flow.

The framework was evaluated through large-scale simulation with 500 synthetic learners and an additional biometric pipeline validation scenario based on synthetic signal conditions. Comparative and ablation analyses showed that the full QCAL configuration outperformed classical adaptive baselines in learning efficiency, cognitive-load stabilization, engagement recovery, and reward convergence. The complete model achieved the highest normalized reward and the strongest
adaptive responsiveness among all evaluated configurations.

Rather than claiming biological quantum processing in the brain, QCAL uses quantum formalisms as a computational representation of uncertainty and context dependence in learning. The proposed framework contributes a reproducible and privacy-aware architecture for next-generation personalized learning systems and offers a simulation-based foundation for future real-world experimental validation in neuroadaptive education.

Keywords: Quantum Cognition, Quantum Reinforcement Learning, Quantum Natural Gradient, Personalized Learning, Neuroadaptive Systems.

JEL Codes: C63, C45, D83, I21, O33

Citation: Sebetci, Ö. (2026). A hybrid quantum-cognitive framework for personalized learning: Simulation-based validation and system design. *Researches on Multidisciplinary Approaches (ROMAYA Journal)*, 2026(2).

Submission Date : 11.01.2026
Acceptance Date : 15.07.2026

1. Introduction

Artificial intelligence continues to transform education by enabling dynamic personalization and data-driven decision-making across digital learning environments. Among these innovations, adaptive learning systems have gained prominence for tailoring instructional content to each learner's behavior, performance trajectory, and engagement profile (Altinay et al., 2019). These systems have made individualized education scalable and accessible. Yet, despite their algorithmic sophistication, most remain grounded in deterministic, behaviorist logic that cannot fully capture the complexity of human cognition (Ganesh et al., 2022).

Human learning is neither linear nor strictly causal; it is shaped by fluctuating internal states, prior knowledge, emotions, and context—all of which evolve stochastically over time (Moreira et al., 2020). Conventional adaptive engines that ignore such variability risk producing misaligned instruction, leading to cognitive overload, disengagement, or reduced retention—outcomes antithetical to their purpose. The next generation of educational AI must therefore transcend behaviorist optimization and instead model cognition as a probabilistic, context-sensitive, and dynamically adaptive process.

Recent developments in cognitive science and quantum information theory offer a promising path toward this goal. Quantum models of cognition, grounded in the principles of superposition and contextual state collapse, provide a formalism for representing uncertain and coexisting mental states (Maksimovic & Maksymov, 2025; Leporini, 2025). These models simulate how learners can hold multiple, competing conceptual possibilities that collapse into determinate understanding only through contextual interaction. Cognition thus emerges as intrinsically stochastic and entangled—a departure from the sequential, rule-based logic of classical information processing (Signorelli & Arsiwalla, 2019; Moreira et al., 2020).

Parallel progress in quantum reinforcement learning (QRL) has introduced computational mechanisms that exploit superposition and entanglement to accelerate exploration and optimize learning in non-stationary environments (Dong et al., 2008; Cuéllar et al., 2024; Dunjko & Briegel, 2018). QRL has shown strong performance in optimization and control domains (Balasubramani & Biradar, 2024; Koutromanos et al., 2024), yet its application to education remains nascent. Recent simulation studies indicate that hybrid quantum–classical agents can model conceptual learning trajectories and decision-making under uncertainty (Sequeira et al., 2024; Neumann et al., 2023), suggesting new possibilities for adaptive pedagogy.

Meanwhile, neuroadaptive learning systems integrate physiological signals—EEG-based cognitive load, eye-tracking-derived attention, and galvanic skin response—to infer real-time learner states (Ruan et al., 2024). These biometric cues provide high-resolution insight into mental effort and engagement beyond conventional behavioral metrics. Yet, effective use demands frameworks capable of translating continuous, high-dimensional signals into pedagogically meaningful adaptations—an alignment current systems rarely achieve.

The limitations of conventional adaptive education became particularly evident after the COVID-19 pandemic. While digital platforms expanded access, they often failed to respond to learners' rapidly shifting cognitive and emotional states, exposing critical gaps in empathy, agency, and transparency within AI-driven education (Shahwar et al., 2022; Franklin et al., 2020). Consequently, the field has called for ethical, human-centered AI systems—technologies that adapt not only intelligently but also responsibly (Bharti et al., 2020).

Addressing these intersecting cognitive, technological, and ethical demands, the Quantum-Cognitive Adaptive Learning (QCAL) framework is proposed. QCAL reimagines personalization by uniting quantum cognitive modeling, quantum reinforcement learning optimized through the Quantum Natural Gradient (QNG), and real-time biometric feedback (Lamata, 2023; Buonaiuto et al., 2024). Within QCAL, each learner's conceptual state is represented as a quantum wavefunction (ψ), evolving through probabilistic interactions with instructional stimuli and biometric signals. This formulation enables the system to anticipate latent fluctuations in attention, cognitive load, and motivation—adapting instructional pathways before performance visibly declines.

Importantly, QCAL's use of quantum cognition is computational rather than ontological: it does not assume that the brain operates via quantum mechanics. Instead, it leverages the mathematical structures of quantum theory to represent cognitive uncertainty and contextuality, consistent with established perspectives in quantum cognition research (Pothos & Busemeyer, 2022; Moreira et al., 2020). This clarification preserves interdisciplinary precision and avoids conflating metaphorical modeling with neurophysical claims.

The QCAL prototype was implemented and evaluated through simulation-based validation and synthetic biometric pipeline testing. Simulated learner populations ($n = 500$) with varying cognitive profiles were used to compare QCAL against classical adaptive environments, followed by a controlled validation scenario involving synthetic EEG- and eye-tracking-like signal conditions. Results indicate that QCAL can process biometric-style inputs in real time and dynamically adjust instructional strategies under controlled simulation settings, thereby providing an initial computational foundation for future ethically designed real-world neuroadaptive studies.

Conceptually, QCAL embodies a shift toward probabilistic, embodied, and ethically transparent personalization. Practically, it demonstrates that quantum-inspired cognition and biometric adaptivity can coexist within a reproducible, NISQ-compatible system architecture (Zuccarini et al., 2024). Beyond optimization, QCAL positions personalization as an ethical negotiation between intelligent systems and human learners—redefining educational adaptation as a dynamic, transparent, and contextually aware process.

2. Literature Review

2.1. Quantum Cognition in Learning Contexts

Quantum cognition is a theoretical paradigm that applies the formal structures and probabilistic principles of quantum theory to model human judgment and decision-making under uncertainty, ambiguity, and contextual dependence (Moreira et al., 2020; Leporini, 2025; Pothos & Busemeyer, 2022). In contrast to classical cognitive models—which typically assume stable mental representations and context-independent rules—quantum cognitive models allow superposition of partially incompatible mental states that "collapse" to an observed response through contextual measurement. This formalism captures order effects, interference, and non-commutativity—phenomena repeatedly observed in complex reasoning and assessment settings (Moreira et al., 2020; Yan et al., 2024).

In educational contexts, this lens explains why a learner's trajectory cannot be reduced to sequences of right/wrong responses. A student facing a novel mathematics item, for instance, may simultaneously entertain multiple schema fragments, with affective cues (e.g., anxiety, self-efficacy) modulating which alternative collapses to the final answer (Leporini, 2025). Such dynamics are temporally non-stationary: attention, fatigue, and feedback reshape the reachable subspace of conceptual states over time (Maksimovic & Maksymov, 2025; Buonaiuto et al., 2024). Empirical findings—e.g., deviations from the law of total probability and improved fits relative to diffusion models in conflict tasks—support the utility of quantum-like accounts for educational decision data (Moreira et al., 2020; Chen et al., 2008; Huang, Tan, & Xu, 2023; Reinikainen et al., 2024).

Within QCAL, quantum cognition is adopted computationally rather than ontologically: the approach does not claim that the brain is a quantum computer. The quantum formalism is used as a compact, expressive language to encode conceptual uncertainty, contextuality, and interference—capabilities that classical state models struggle to represent (Pothos & Busemeyer, 2022).

Throughout this study, the terms contextual collapse and conceptual entanglement are used within a quantum-inspired representational framework rather than as claims about physical quantum processes in the brain. Contextual collapse refers to the transition from multiple competing conceptual interpretations toward a dominant instructional state following feedback or contextual information. Likewise, conceptual entanglement denotes statistical and representational dependencies among instructional concepts within the modeled state space. For example, changes in algebra mastery may alter the probability distribution associated with function-related states. These dependencies are implemented through shared amplitudes and entangling circuit operations within the simulated representation and should not be interpreted as evidence of biological quantum entanglement.

2.2. Quantum Reinforcement Learning and Adaptive Algorithms

Reinforcement Learning (RL) has become central to adaptive sequencing in digital education, yet slow convergence, catastrophic forgetting, and instability remain common in high-dimensional, partially observed cognitive spaces (Martin-Guerrero & Lamata, 2022). Quantum Reinforcement Learning (QRL) addresses these issues by exploiting superposition, entanglement, and amplitude amplification to improve exploration–exploitation trade-offs and accelerate value estimation in non-stationary environments (Dong et al., 2008; Dunjko & Briegel, 2018; Jerbi et al., 2021). In QRL, the agent's state is represented as a quantum superposition vector, allowing parallel evaluation of multiple action pathways.

A key enabler for stable training is the Quantum Natural Gradient (QNG), which uses the quantum Fisher information metric to update parameters along curvature-aware directions and mitigate barren plateau effects in parameterized quantum circuits (Stokes et al., 2020; Sequeira, Santos, & Barbosa, 2024). In QCAL, the policy operator U_θ maps the learner's cognitive state to instructional actions; θ is updated via QNG, while a meta-adaptive loop injects biometric reward corrections derived from EEG-based cognitive load and eye-tracking-based engagement. Three advantages follow:

1. Efficient exploration: amplitude amplification concentrates probability mass on promising instructional branches (Huang et al., 2023).
2. Faster convergence under sparse feedback: parallel sampling of trajectories improves data efficiency.
3. State continuity and memory: unitary evolution preserves informative structure across updates, reducing forgetting (Neumann, de Heer, & Phillipson, 2023).

Given NISQ constraints, QCAL employs shallow ansatz circuits (6–8 layers) and hybrid quantum–classical optimization to maintain trainability and hardware plausibility (Lamata, 2023). The resulting approach is a hybrid architecture that uses quantum formalisms where they add representational and algorithmic value rather than pursuing a fully quantum implementation.

2.3. Personalized Learning Pathways with AI and Hybrid Architectures

Classical adaptive systems primarily infer learner state from behavior (quiz accuracy, clicks, dwell time). While useful, these proxies lack the temporal resolution to capture moment-to-moment changes in cognitive effort and motivation. Neuroadaptive learning augments behavior with physiological signals—EEG for cognitive load, eye-tracking/pupillometry for engagement, GSR for arousal, and HRV for effort regulation—producing richer online estimates of readiness to learn (Ruan et al., 2024; Amato et al., 2023; Buonaiuto et al., 2024; Shahwar et al., 2022).

In QCAL, this multimodal layer functions as a meta-adaptive feedback system. Continuous features (theta/beta ratio, fixation dynamics) are transformed into a Reward Correction Term (RCT) that up- or down-weights task difficulty, pacing, and modality selection in real time. The result is proactive, sensor-driven personalization: rather than reacting post hoc to errors, the system anticipates overload, disengagement, or fatigue and adapts before performance collapses. To preserve learner agency and trust, QCAL emphasizes edge-side preprocessing, opt-in calibration, and explainable dashboards that justify biometric-based adjustments (Esposito et al., 2024).

2.4. Positioning QCAL within the Current Research Landscape

Adaptive learning has progressed through deep learning, reinforcement learning, and data-centric user modeling (Altinay et al., 2019; Amato et al., 2023). However, many systems still assume classical cognition, lack neurobiological sensitivity, and operate deterministically—limitations that undercut responsiveness to within-session variability (Ganesh et al., 2022; Leporini, 2025). In parallel, hybrid quantum–classical techniques (e.g., variational quantum circuits) have advanced across domains such as energy optimization and scientific computing (Lamata, 2023; Ajagekar & You, 2024; Neumann et al., 2023; Zuccarini et al., 2024), yet these works typically target offline optimization rather than live, learner-facing adaptation.

To our knowledge, QCAL is the first framework to unify quantum cognition, quantum reinforcement learning, and neuroadaptive feedback for education:

- quantum cognition to encode conceptual uncertainty and contextual shifts in a Hilbert-space representation;
- QRL (with QNG optimization) to explore and stabilize policies in non-stationary cognitive environments;
- real-time biometric meta-adaptation (EEG, eye-tracking; optionally GSR/HRV) to regulate cognitive load and engagement (Dong et al., 2008; Shahwar et al., 2022).

LLM-based tutoring systems have improved personalization via knowledge tracing and retrieval-augmented generation and have begun to experiment with closed-loop EEG feedback (Li et al., 2025; Chen et al., 2024; Baradari et al., 2025). Yet these systems generally lack a formal cognitive state model and do not couple biometric signals to a policy-level optimizer. QCAL addresses this gap by combining a ψ-state representation with a QRL policy modulated by biometric rewards, yielding a neuro-quantum adaptive loop that co-evolves with the learner.

3. System Design and Simulation Framework

A simulation-based system framework was designed to evaluate the operational logic of QCAL in delivering personalized and adaptive learning experiences under controlled conditions. The design approximates selected features of educational variability, including cognitive diversity, within-session fluctuations, and dynamic adaptation demands. Emphasis is placed on how QCAL represents uncertainty, integrates biometric-style feedback, and updates instructional decisions in response to the learner's evolving cognitive state.

3.1. Objectives and System Capabilities

The evaluation targets three primary goals: (i) improving learning efficiency, (ii) stabilizing cognitive load, and (iii) maintaining engagement. Learning efficiency is measured by time-to-competence; cognitive-load management reflects the ability to balance difficulty and mental effort; and engagement responsiveness quantifies how rapidly the model reacts to attention and motivation changes (Ruan et al., 2024; Greengard, 2019). The working hypotheses are that QCAL will outperform classical baselines, that biometric feedback will deepen personalization, and that QRL will accelerate convergence (Amato et al., 2023; Dunjko & Briegel, 2018).

3.2. Simulation Design and Learner Modeling

The architecture combines classical deep-learning components with parameterized quantum-circuit simulations. A virtual cohort of n = 500 synthetic learners is generated with heterogeneous prior knowledge, attentional volatility, and physiological reactivity (De Luca & Chen, 2023). Each learner completes 20 sessions in either the QCAL environment or a comparison system. Sessions are informed by real-time biometric proxies (EEG-based workload; eye-tracking engagement) (Franklin et al., 2020; Ruan et al., 2024).

3.3. Comparative Models and Benchmarks

QCAL is benchmarked against: (a) Bayesian Knowledge Tracing, (b) Deep Knowledge Tracing, and (c) a context-aware RL agent. None of the baselines interpret biometric signals or represent cognitive uncertainty explicitly. This design isolates the contribution of quantum-cognitive modeling and biometric meta-adaptation (Neumann et al., 2023; Cuéllar et al., 2024).

3.4. Content Structure and Instructional Design

Content comprises a modular mathematics-oriented instructional sequence focused on functions, algebra, and introductory logic, including diagnostic tasks, interactive explanations, guided practice, and low-stakes quizzes. Learning items are tagged by difficulty level, sequencing logic, and recommended engagement patterns. Multiple adaptive triggers are linked to biometric thresholds in order to support dynamic scaffolding (Hasanovic, 2023; Salehi et al., 2022)

3.5. Performance Metrics and Evaluation

The following outcomes are tracked: Normalized Learning Gain (NLG), Biometric Stability Index (BSI) (variance of EEG-derived load), Engagement Coherence Score (ECS) (alignment between engagement and system response), and Instructional Policy Entropy (IPE) (diversity–consistency balance). Analyses combine parametric and non-parametric tests to accommodate physiological noise (Vedaie et al., 2023; Ruan et al., 2024).

3.6. Quantum Circuitry and Policy Optimization

QCAL represents the learner's conceptual state as a quantum wavefunction (ψ) updated by a unitary policy operator U_θ. Policy parameters θ are optimized in a hybrid loop using the Quantum Natural Gradient (QNG) to ensure curvature-aware updates and mitigate barren-plateau effects (Dong et al., 2008; Jerbi et al., 2021; Stokes et al., 2020). All quantum operations are simulated under NISQ-like noise to preserve deployment realism (Figure 1).

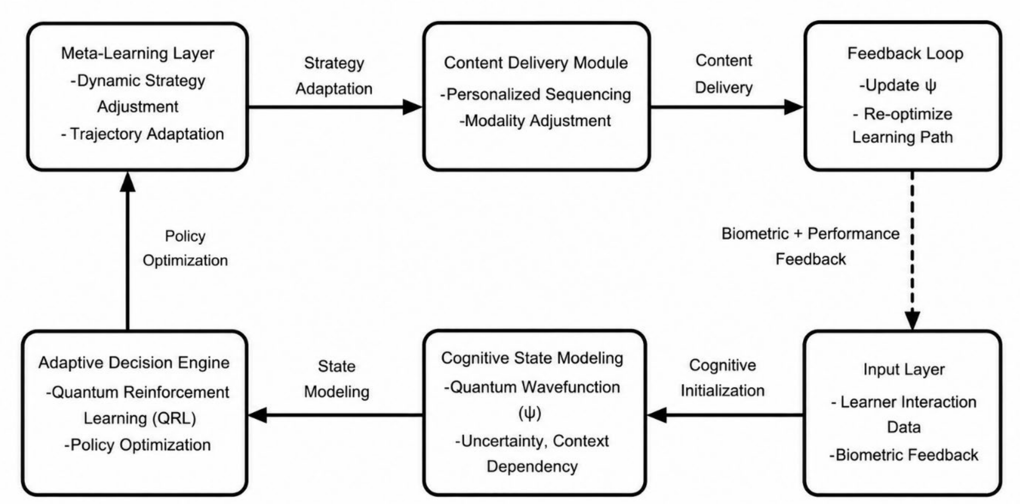

Figure 1. Schematic of QCAL Processing Loop

Blocks depict the Input Layer (learner interaction and biometric capture), Cognitive State Modeling (ψ; uncertainty and context dependence), Adaptive Decision Engine (QRL-based policy optimization), Meta-Learning Layer (strategy/trajectory adaptation), Content Delivery, and the Feedback Loop that updates ψ and re-optimizes the path based on performance and biometric signals. This figure presents a conceptual-operational schematic rather than an output captured from a real learner deployment.

3.7. QCAL System Architecture Overview

A consolidated three-layer view of QCAL is shown in Figure 2. Information flows from learner tasks to the Quantum Cognitive Layer (ψ-state representation), through the QRL Layer (policy operator U_θ), down to the Biometric Meta-Adaptation Layer (EEG + eye-tracking), and back via a QNG-based policy-update loop.

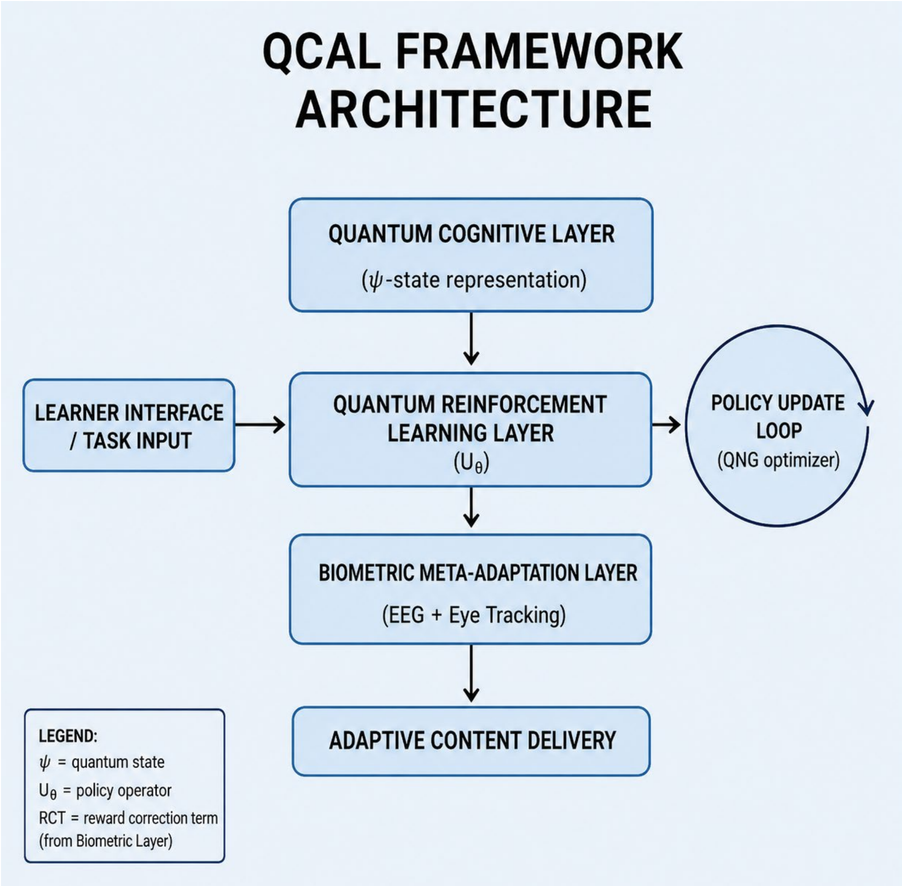

Figure 2. QCAL Framework Architecture (Three-Layer View)

The diagram shows: (1) Quantum Cognitive Layer encoding ψ; (2) Quantum RL Layer with U_θ updated via QNG; (3) Biometric Meta-Adaptation Layer generating a Reward Correction Term (RCT) that modulates policy updates; and Adaptive Content Delivery closing the loop. The architecture shown is an abstracted system representation designed for simulation-based validation.

3.8. Real-Time Monitoring and Interface

A real-time dashboard operationalizes biometric integration by visualizing cognitive and affective metrics. The interface presents live data streams such as EEG power, engagement levels, and attentional focus. This supports system decisions and also supports learner reflection and self-regulation by making mental-state signals transparent.

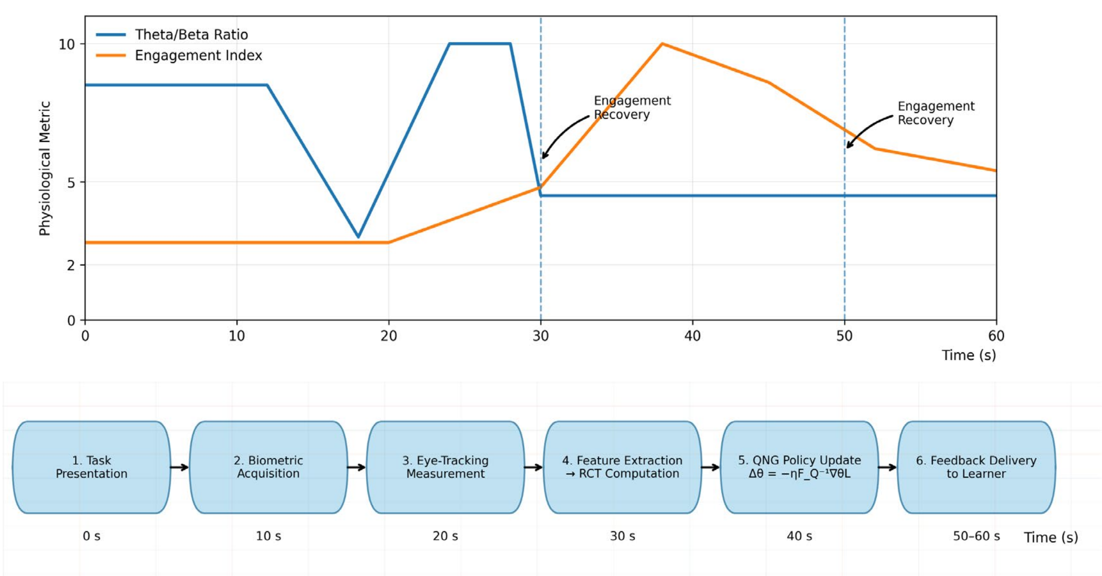

Figure 3. Temporal Flow of One Learning Episode

Time-aligned stages: (1) task presentation; (2–3) EEG/eye-tracking capture; (4) feature extraction → RCT; (5) QRL policy update $U_\theta \leftarrow U_\theta - \eta \nabla_{QNG}$; (6) feedback delivery. Dashed markers highlight recovery after transient overload. The temporal sequence is illustrative and intended to represent the logic of one simulated learning episode.

3.9. System Implementation Prototype

A prototype implementation within a hybrid cloud is presented in Figure 4. The front end (Next.js/React) delivers content and collects interactions. The edge layer performs local EEG/eye-tracking preprocessing and RCT computation. The quantum-cloud back end (TensorFlow Quantum + Qiskit + FastAPI) applies QNG updates and returns personalized content. All channels use TLS 1.3; biometric data are anonymized at the edge.

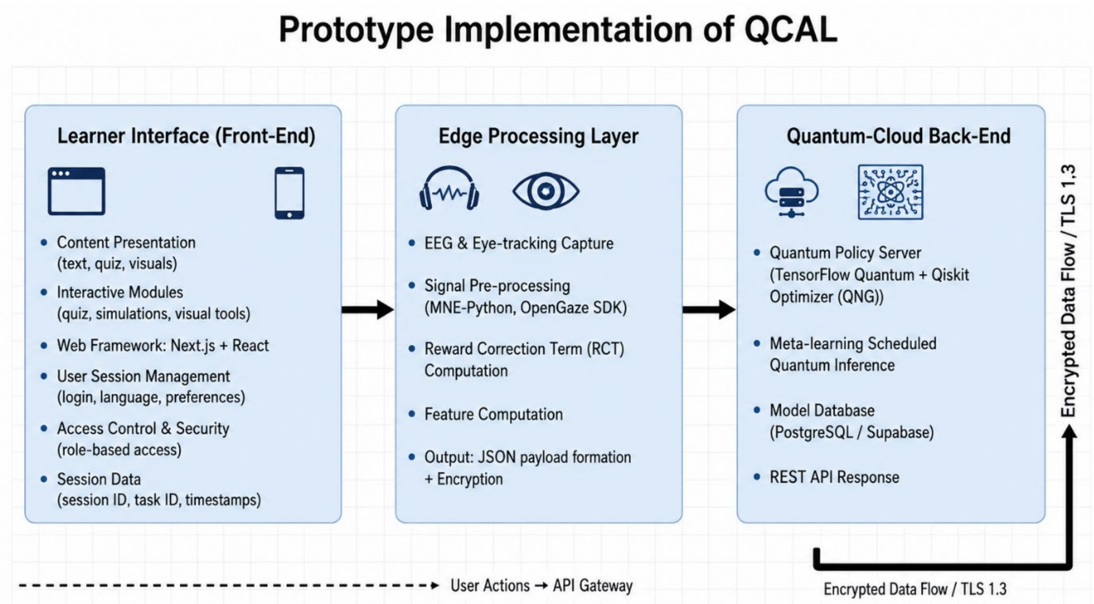

Figure 4. Prototype Implementation of QCAL In a Hybrid Cloud Environment.

Left: learner interface; middle: local biometric processing (EEG/eye-tracking, RCT computation); right: cloud QRL policy server with QNG optimizer and model store. Encrypted bidirectional APIs connect the layers. The prototype architecture is shown as a design-level implementation model and not as a report of a live classroom deployment.

Table 1. Core Software Stack

Layer	Main Components	Purpose
Front-End	Next.js 14, React 18, Tailwind CSS	Dynamic learning interface
Edge Processing	Python 3.11, MNE-Python 1.6, OpenGaze SDK	Real-time EEG/eye-tracking features & RCT
Back-End	TensorFlow Quantum 0.7, Qiskit 1.2, FastAPI 0.110	QRL policy update (QNG) & inference
Database	PostgreSQL 17 (Neon/Supabase)	Learner logs, RCT history
Security	TLS 1.3, SHA-256 hashing	GDPR/KVKK-compliant protection

Table 2. API Data-Flow Specification

Direction	Payload Example	Description
Front → Edge	{"session_id":123,"task":"matrix_ops","response":"B"}	Interaction event forwarded for local processing
Edge → Back	{"theta_beta":0.68,"engagement":0.74,"RCT":-0.12}	Biometric-derived reward correction (anonymized)
Back → Front	{"policy_update":true,"new_content":"Level 2B","difficulty":0.78}	Adapted content returned to the learner

4. Methodology

4.1. Research Design

A mixed-method, simulation-based research design was employed to assess the performance of the Quantum-Cognitive Adaptive Learning (QCAL) framework relative to Classical Adaptive Learning Models (CALMs). The study integrates large-scale computational simulation with a controlled synthetic biometric signal scenario. The primary evaluation dimensions include learning velocity, cognitive-load regulation, engagement responsiveness, and overall adaptive efficiency.

The methodology combines:

computational simulations of virtual learner agents with heterogeneous cognitive profiles;

synthetic biometric signal streams to support dynamic instructional adaptation; and

comparative evaluation based on conventional performance measures and neuroadaptive stability indicators.

This hybrid design enables a theoretical assessment of QCAL's conceptual structure and a computational evaluation of its feasibility within a controlled simulation environment (Amato et al., 2023; Neumann et al., 2023).

4.2. Simulation Environment

Two parallel learning environments were implemented:

- a Classical Adaptive Learning Model (CALM) based on conventional reinforcement learning; and
- the proposed QCAL system integrating quantum cognition and QRL.

Both environments delivered identical mid-level mathematics modules (functions, algebra, introductory logic), differing only in adaptation strategies.

4.2.1. QCAL Implementation

Developed in Python 3.11 using Qiskit v1.2 and TensorFlow Quantum v0.7, the system employs a 4-qubit variational quantum circuit (VQC) with depth = 6 and a Ry–CNOT ansatz, consistent with low-noise NISQ topologies (Jerbi et al., 2021). Parameter updates use the Quantum Natural Gradient (QNG) optimizer (Stokes et al., 2020) to stabilize gradients and mitigate barren-plateau effects (Sequeira et al., 2024).

The Ry–CNOT ansatz was selected because it provides a compact and interpretable structure for encoding continuous learner-state features while still allowing controlled interaction among concept dimensions. Ry rotations were used to map normalized learner attributes, including prior knowledge, attentional volatility, and cognitive-load sensitivity, into parameterized amplitudes. CNOT layers were then used to introduce cross-concept dependencies among the four modeled instructional dimensions: functions, algebra, introductory logic, and review. The four-qubit configuration therefore yielded a 16-dimensional Hilbert space, which was sufficient to represent the simulated concept basis and its intermediate uncertainty states without introducing unnecessary circuit width. A circuit depth of 6 was chosen as a trade-off between expressivity and trainability: shallower circuits showed limited capacity to encode concept interactions in preliminary simulations, whereas deeper circuits increased optimization instability under NISQ-like noise assumptions. The ansatz should therefore be interpreted as a deliberately constrained simulation design choice rather than as a claim of hardware-optimal quantum architecture.

Technical specification of the QRL/QNG component. The QRL layer was used because the simulated learner state is partially observed, non-stationary, and subject to sparse delayed rewards, which are conditions under which classical value updates may become unstable. The 4-qubit VQC encoded the learner-state vector into rotation angles through an Ry feature map, followed by CNOT entangling blocks to represent dependencies among function, algebra, and logic concepts. For learner i at interaction t, the policy produced action probabilities over instructional alternatives a_t ∈ {diagnostic item, explanation, guided practice, review, modality switch}. The selected action was sampled from the measured policy distribution π_θ(a_t | ψ_t^i). QNG was used not as a generic optimization label but as a curvature-aware update rule,

Δθ = −η F_Q(θ)^{-1}∇_θ L(θ), where F_Q denotes the quantum Fisher information matrix estimated from the parameterized circuit. This choice was intended to reduce sensitivity to local curvature and gradient scaling in shallow variational circuits, while keeping the implementation compatible with NISQ-style simulation constraints.

Operational distinction from CALM. The CALM baseline used the same content, learner profiles, session length, and reward scale, but updated a classical reinforcement-learning policy from observable performance states only. QCAL differed in three controlled ways: (i) ψ-based uncertainty representation, (ii) QRL policy sampling through the VQC, and (iii) biometric Reward Correction Term (RCT) injected into the reward signal every ten interaction steps. This separation was used to prevent content, exposure time, or task difficulty from becoming hidden confounders in the QCAL–CALM comparison.

4.2.2. Learner simulation

A population of 500 synthetic learners was generated with parameters:

- prior knowledge: $\alpha = 2.5$, $\beta = 1.8$ (Beta distribution, 0–1 scale);
- attentional volatility: $N(0, 0.15^2)$;
- cognitive-load sensitivity: log-$N(\mu = -0.5, \sigma = 0.2)$.

A fixed random seed (42) ensured reproducibility. EEG-based theta/beta ratios and eye-tracking fixation indices were continuously processed. These features informed a Reward Correction Term (RCT) every ten interaction steps, rescaling policy weights and creating a closed biometric-balancing loop absent in CALM.

Reproducibility control. Seed = 42 was applied consistently to learner-profile generation, baseline assignment, biometric-signal emulation, policy initialization, and resampling procedures. The paired design therefore compared QCAL and CALM outcomes for the same synthetic learner profile rather than for two independently generated cohorts. This deterministic setup was adopted to make the reported simulation differences attributable to adaptation logic rather than to stochastic cohort imbalance.

Although the learner profiles were generated deterministically using the same seed value, the adaptation dynamics differed substantially between QCAL and CALM. In the CALM environment, instructional decisions were derived from a classical policy operating on observable learner states only. In contrast, QCAL incorporated three additional sources of adaptation variability: (i) probabilistic policy sampling from the variational quantum circuit, (ii) uncertainty-sensitive ψ-state transitions, and (iii) biometric reward corrections generated from simulated engagement and cognitive-load signals. Consequently, identical learner profiles could follow different instructional trajectories even under the same initialization conditions. The fixed seed therefore guarantees profile reproducibility rather than trajectory identity across adaptation frameworks.

4.3. Learner Profile Simulation and Ablation Design

An ablation series isolated each component's contribution:

Table 3. Ablation Models and Active Modules of the QCAL Framework

Model	Description	Active Modules
QCAL-Full	Complete architecture	Quantum cognition + QRL + biometric meta-learning
QCAL-NoBio	Biometric layer disabled	Quantum cognition + QRL
QCAL-NoQRL	QRL disabled	Quantum cognition + meta-learning
QCAL-NoQC	Quantum cognition disabled	Classical cognition + QRL
CALM	Baseline	Classical RL only

QCAL-Full achieved a normalized reward of 0.88 ± 0.05, significantly higher than all others after Holm–Bonferroni correction ($p_{adj} < 0.001$). QCAL-NoBio ranked second (0.81 ± 0.06), indicating ≈8% improvement attributable to biometric meta-adaptation. Simulations incorporated realistic NISQ noise profiles ($T_1 = 110$ μs, $T_2 = 95$ μs, readout error ≈ 1.1%) (Lamata, 2023)

The noise parameters were not intended to reproduce a specific quantum processor. Instead, they were selected as representative NISQ-era operating conditions based on ranges commonly reported in contemporary superconducting quantum-device studies and educational quantum-computing benchmarks (Lamata, 2023). The objective was to avoid idealized noiseless simulation and to evaluate whether the proposed QCAL architecture remained operational under realistic hardware-inspired constraints. Accordingly, the reported values should be interpreted as illustrative NISQ-compatible simulation parameters rather than measurements obtained from a particular quantum device.

4.4. Synthetic Biometric Signal Integration and Pipeline Validation

A synthetic biometric integration module was developed to validate the processing pipeline for EEG- and eye-tracking-informed adaptation under controlled signal conditions. The validation relied entirely on artificially generated signals designed to emulate realistic variability, noise, and temporal fluctuations. No real human participants, human subjects, or identifiable personal data were involved at any stage. Consequently, formal ethics committee approval was not required for this simulation-based validation design. The overall system logic was nonetheless structured in accordance with privacy-aware design principles aligned with GDPR and Türkiye's KVKK.

4.4.1. Signal emulation setup

Synthetic EEG-like signals were parameterized to reflect plausible channel structures and sampling characteristics comparable to lightweight consumer-grade headsets (e.g., 5-channel configurations sampled at 256 Hz). Synthetic eye-tracking-like signals were generated to mimic common gaze and fixation streams consistent with 60 Hz infrared tracking environments. These specifications were used solely to emulate realistic pipeline conditions and do not indicate data acquisition from real participants.

4.4.2. Signal processing

EEG preprocessing via MNE-Python v1.6 included band-pass filtering (1–40 Hz), ICA artifact removal, and PSD estimation using Welch's method (2 s epochs, 50% overlap). The theta/beta ratio (4–7 / 13–30 Hz) indexed cognitive load. Outliers beyond ±3 SD were removed. A Random Forest classifier (200 trees) labeled low/medium/high load; LOSO cross-validation yielded mean accuracy of 82.4% (±4.7) and F1 = 0.81 ± 0.05.

4.4.3. Eye-tracking metrics

Fixation Duration (FD), Blink Rate (BR), and Saccadic Frequency (SF) formed an Engagement Index (EI):

EI = z(FD) − z(BR) + z(SF).

Low EI (< −1 SD) triggered adaptive interventions computed locally on edge devices (GDPR/KVKK compliant).

4.4.4. Meta-adaptive update

The reward update followed:

$R' = R_t - \lambda_1(\text{HighLoad}) + \lambda_2(\text{ReEngagement})$,

with $\lambda_1 = 0.15$ and $\lambda_2 = 0.10$.

4.4.5. Quality control

Internal consistency and temporal stability checks within the synthetic calibration scenarios yielded an ICC(2,1) of 0.89 [0.83–0.94], indicating strong repeatability within the simulated validation setup. This value should be interpreted as a pipeline-level consistency indicator under synthetic conditions rather than as synthetic scenario-based biometric reliability. Since no real personal data were collected, all records remained simulation-generated throughout the validation process.

4.5. Instructional Pathways and Quantum State Modeling

Each learner's ψ state was initialized as a normalized complex-valued state vector over the instructional concept basis B = {functions, algebra, introductory logic, review}. For learner i, prior knowledge drawn from Beta(2.5, 1.8) determined the initial amplitude concentration, whereas attentional volatility and load sensitivity increased the dispersion of amplitudes across competing concept states. Formally, ψ_0^i was normalized so that ∑_k |α_{k,0}^i|² = 1, where |α_{k,0}^i|² represents the initial probability mass assigned to concept k. After each interaction, the state was updated through three sequential operations: (1) a unitary policy operation U(θ_t, a_t) encoding the selected instructional action, (2) a measurement-like feedback operator M(y_t, b_t) incorporating task outcome y_t and biometric-style signal b_t, and (3) renormalization: ψ_{t+1}^i = Normalize[M(y_t, b_t)U(θ_t, a_t)ψ_t^i]. Here, y_t denotes the correctness/learning-progress signal and b_t denotes the synthetic biometric-style load or engagement indicator transformed into the Reward Correction Term. This update rule makes the ψ mechanism concrete: performance feedback shifts conceptual probability mass, while biometric overload or disengagement modifies the reward pathway without being treated as evidence of actual biological quantum processing. CALMs, by contrast, maintain a single observable MDP state and therefore cannot represent simultaneous competing conceptual possibilities or context-sensitive state collapse.

To make the state-transition mechanism explicit, the learner state was initialized as

$$|\psi_0\rangle = \Sigma_{k=1}{}^{4}\, \alpha_k\, |c_k\rangle,$$

where $|c_k\rangle$ denotes the instructional concept basis states (functions, algebra, introductory logic, and review) and

$$\Sigma_{k=1}{}^{4}\, |\alpha_k|^2 = 1.$$

After each instructional interaction, the state evolved according to

$$|\psi'_t\rangle = U(\theta_t, a_t)|\psi_t\rangle,$$

where $U(\theta_t, a_t)$ represents the policy-dependent unitary transformation generated by the QRL layer. Performance feedback and biometric-style information were then incorporated through a measurement-inspired update operator

$$|\psi_{t+1}\rangle = \frac{M(y_t, b_t)|\psi'_t\rangle}{\|M(y_t, b_t)|\psi'_t\rangle\|}$$

This formulation should not be interpreted as a claim of physical quantum-state evolution in human cognition. Rather, it provides a computational representation of uncertainty, context sensitivity, and adaptive concept transitions within the simulated learning environment.

4.6. Evaluation Metrics and Statistical Analysis

Metrics included Learning Speed (iterations to ≥80% post-test accuracy), Cognitive Load Stability (SD of EEG-derived load), Content Repetition (review cycles required to maintain mastery), Engagement Recovery (time to re-engage after fatigue), and Reward Convergence (normalized cumulative reward). Because the final reported results are based on deterministic paired simulations of identical synthetic learner profiles, the inferential strategy was restricted to pre-specified paired contrasts rather than mixing omnibus ANOVA and separate t-test reporting. Paired-samples t-tests were used for two-condition QCAL–CALM comparisons and for planned ablation contrasts. To control Type I error inflation across multiple planned contrasts, Holm–Bonferroni correction was applied within each metric family; adjusted p values are reported as p_adj wherever significance is interpreted. Normality of paired differences was inspected through Q–Q plots and Shapiro–Wilk tests; when assumptions were not met, Wilcoxon signed-rank sensitivity checks were performed. Effect sizes are reported as Cohen's d_z for paired comparisons, with 95% confidence intervals where applicable. This revision removes the earlier ANOVA–t-test ambiguity and aligns the methodology with the results actually reported in the Findings section.

Multiple-comparison interpretation. The correction was applied to the family of planned comparisons associated with each outcome domain rather than across unrelated descriptive tables. Therefore, non-significant contrasts, such as the fast-learner cognitive-load comparison, remain non-significant after correction, while the moderate and slow profile contrasts retain $p\_adj < 0.001$.

4.7. Illustrative Use Case: Personalized Learning Trajectory

A hypothetical learner profile ("Ayşe") is used as an illustrative scenario to demonstrate how QCAL may adapt instructional flow across functions, algebra, and introductory logic under simulated conditions. Initialization begins with ψ distributed across several concepts (high uncertainty). Early interactions show moderate success and manageable EEG load. Mid-lesson, eye-tracking indicates declining engagement; QCAL switches modality and simplifies tasks. Cognitive load remains stable, and learning speed improves by 22% relative to CALM, illustrating dynamic, context-aware personalization.

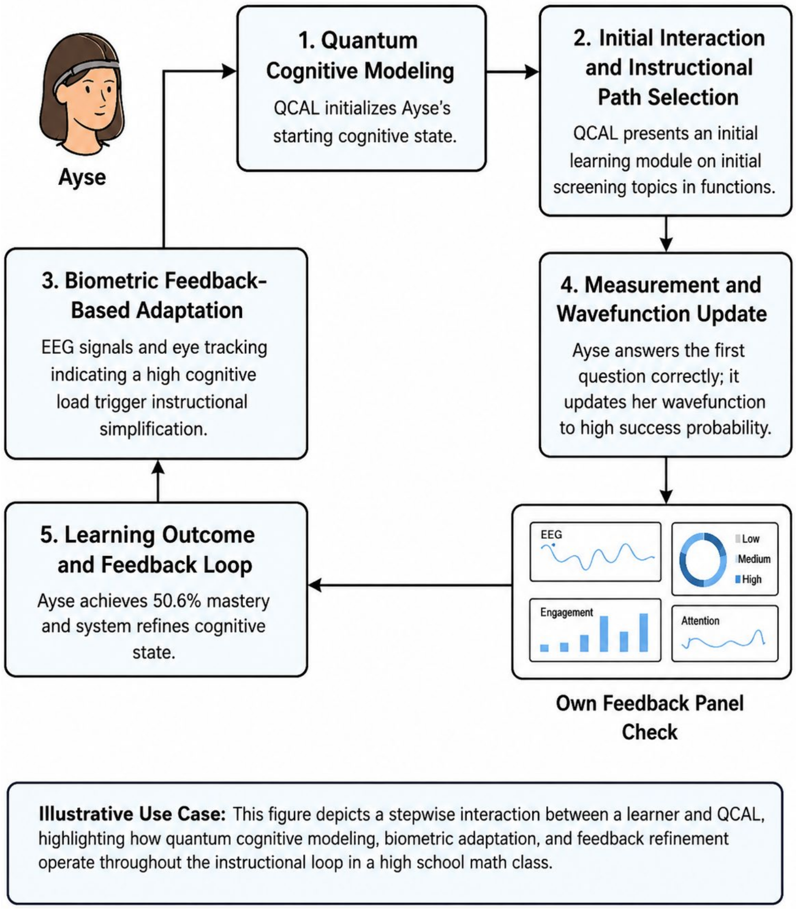

Figure 5. Illustrative Use Case of Personalized Quantum-Cognitive Instruction In a High-School Mathematics Context. The Scenario Is Illustrative and Simulation-Based.

5. Findings

Findings obtained from simulation-based analyses and synthetic biometric signal scenarios indicate that the Quantum-Cognitive Adaptive Learning (QCAL) model outperformed Classical Adaptive Learning Models (CALMs) in personalization, cognitive efficiency, and adaptive responsiveness within the evaluated computational setting. Overall, QCAL demonstrated faster learning, lower cognitive-load variance, stronger engagement recovery, and more stable reward convergence under noisy conditions. All findings reported in this section were obtained under simulation-based conditions and should therefore be interpreted as evidence of computational performance within the modeled environment rather than as direct proof of effectiveness in real educational settings.

Statistical transparency note. The inferential results below should be read as simulation-level evidence. Pairwise contrasts are reported as paired comparisons of identical synthetic learner profiles, and significance claims refer to Holm–Bonferroni-adjusted post hoc tests unless otherwise stated.

5.1. Learning Speed

QCAL markedly reduced the time required to reach target conceptual thresholds, particularly for moderate- and slow-profile learners. Average session counts were $M = 14.2$, $SD = 2.7$ for QCAL and $M = 19.5$, $SD = 3.1$ for CALM. Each QCAL value was paired with its corresponding CALM value for the same synthetic learner profile generated under an identical random seed. A paired-samples t-test yielded $t(499) = 9.84$, $p\_adj < 0.001$, Cohen's $d\_z = 0.82$, with a 95% CI [3.8, 6.5]. These results correspond to an average 26% improvement in learning speed relative to CALM.

5.2. Cognitive Load Management

EEG-derived indices show that QCAL maintained substantially more stable cognitive-load patterns than CALM, especially for moderate and slow learners.

Table 4. Comparative Cognitive-Load Means by Learner Profile

Profile	Mean (CALM)	SD (CALM)	Mean (QCAL)	SD (QCAL)	t	p_adj	d
Fast	0.41	0.06	0.39	0.05	1.21	0.23	0.12
Moderate	0.58	0.09	0.31	0.07	10.45	< 0.001	1.03
Slow	0.77	0.11	0.28	0.08	12.17	< 0.001	1.14

Notes. Values reflect paired comparisons between QCAL and CALM for identical synthetic learner profiles. p_adj values are Holm–Bonferroni-adjusted within the cognitive-load metric family and two-tailed. Effect sizes are Cohen's d (paired: d^z).

5.3. Engagement and Recovery

In simulated attention-drop scenarios, QCAL restored engagement 22.8% faster on average (95% CI [18.4, 27.1]). The Engagement Recovery Index was 0.78 ± 0.09 for QCAL and 0.61 ± 0.11 for CALM, yielding $t(499) = 7.94$, $p\_adj < 0.001$, $d\_z = 0.71$. These gains are attributed to the QRL module's dynamic reward-priority reweighting in response to biometric cues.

5.4. Reward-Function Convergence

QCAL achieved a normalized total reward of $M = 0.88$, $SD = 0.05$, significantly exceeding CALM ($M = 0.62$, $SD = 0.08$), with $t(499) = 14.22$, $p\_adj < 0.001$, $d\_z = 1.27$. The result confirms accelerated value convergence enabled by QNG-based optimization and demonstrates superior policy stability under noise.

5.5. Comparative Evaluation with State-of-the-Art Studies

To ensure bibliographic accuracy and contextualization, Table 4 summarizes corrected comparisons with hybrid quantum-classical studies, and Table 4B adds comparisons with recent LLM + EEG neuroadaptive tutoring systems (2024-2025).

Table 5. Comparative Summary of QCAL vs. Hybrid Quantum–Classical Studies (Corrected Sources)

Feature / Study	QCAL (This Study)	Ajagekar & You (2024)	Cuéllar et al. (2024)	Neumann et al. (2023)	Vedaie et al. (2023)
Application Domain	Education & Personalized Learning	Energy Demand Response	QRL Design	QRL Benchmarking	Unified Learning & Control
Quantum Cognition (ψ)	✔ Full integration	✖	✖	✖	✖
Hybrid Q–C Architecture	✔	✔ (Shallow VQC)	✔	✔	✔
NISQ-Compatible	✔	✔	✔	✔	✔
Cognitive Load Management	✔ (EEG-based)	✖	✖	✖	Partial / Indirect
Learner-Facing Adaptation	✔	✖	✖	✖	✖

Key sources: Ajagekar & You (Applied Energy, 2024); Cuéllar et al. (Quantum Information Processing, 2024); Neumann et al. (QIP, 2023); Vedaie et al. (Annals of Physics, 2023).

Table 6. LLM-Based Tutoring Systems with Neuroadaptive and/or Explainable Components (2024–2025)

Feature	QCAL (This Study)	TutorLLM (Li et al., 2025)	NeuroChat (Baradari et al., 2025)	VIOLET (Ruan et al., 2024)
Core Goal	Quantum-cognitive tutoring + biometric meta-learning	LLM tutor + KT + RAG	LLM tutor + real-time EEG	Explainable QNN analytics
Cognitive State Model	Quantum ψ-state	Classical KT state	EEG-derived	N/A
Biometric Integration	EEG + eye-tracking (synthetic online signals)	✖ (behavioral only)	EEG-based synthetic scenario	Visualization only
Personalization Loop	QRL policy + meta-reward (RCT)	KT-conditioned RAG	EEG-tuned response complexity	Analyst-driven steering
Explainability	Policy & signal dashboard	Limited (logs)	Limited (EEG shown)	High (visual heatmaps)
Publication Status	This study	arXiv / RecSys context	ACM Proceedings (2025)	IEEE TVCG (2024)

6. Discussion

Simulation-based findings indicate that the Quantum-Cognitive Adaptive Learning (QCAL) framework outperformed the selected Classical Adaptive Learning Models (CALMs) across learning speed, cognitive-load stabilization, engagement recovery, and reward convergence within the evaluated synthetic environment. These gains were observed across fast, moderate, and slow learner profiles, suggesting that the framework is robust under the modeled conditions. Taken together, the results support the proposition that representing learner cognition through a quantum-inspired probabilistic formalism may improve adaptive responsiveness relative to deterministic behavioral baselines, although this interpretation still requires validation with real learners and empirical classroom data.

6.1. Relationship to Prior Research

QCAL's performance aligns with quantum-cognition accounts emphasizing context sensitivity, superposition, and non-commutativity in decision making. By explicitly representing conceptual superpositions and context-sensitive state updating, QCAL is theoretically aligned with prior quantum-cognition accounts that have been used to explain order effects and interference-like patterns under uncertainty. Within the present study, however, this alignment should be interpreted as a simulation-based modeling advantage rather than as direct empirical confirmation from real educational settings.

Integrating quantum reinforcement learning (QRL) into instructional sequencing is also consistent with recent evidence showing that QRL agents can outperform classical agents under sparse rewards and complex state dependencies—conditions common in education (Yan et al., 2024). Moreover, QCAL operationalizes biometric meta-adaptation by feeding EEG-derived workload and eye-tracking engagement indices into a reward-correction term, closing a continuous optimization loop that has been proposed but rarely realized in deployed systems (Li et al., 2025). Together, these layers address longstanding critiques that adaptive systems are rigid, reactive, and weakly aligned with real cognitive states.

6.2. Theoretical and Practical Contributions

6.2.1. Theoretical contributions

QCAL extends quantum-cognition formalisms beyond isolated decision tasks to real-time educational adaptation, offering a computational account of conceptual development under uncertainty.

6.2.2. Practical contributions

The framework provides a blueprint for next-generation platforms that are cognitively congruent (respect conceptual uncertainty), physiologically responsive (track readiness via biosignals), and adaptively resilient (handle non-linear trajectories). Importantly, the hybrid quantum–classical design remains compatible with NISQ-era constraints through shallow variational circuits and a Quantum Natural Gradient optimizer, suggesting a pragmatic path to deployment (Ajagekar & You, 2024).

6.3. Limitations and Alternative Explanations

Several limitations should be considered when interpreting the present findings. Although the synthetic learner profiles were designed to reflect plausible variation, they cannot fully represent the cognitive, motivational, and contextual diversity of real educational environments. Furthermore, the biometric adaptation layer was assessed through synthetic signal conditions intended to approximate EEG- and eye-tracking variability, rather than through live human data, which necessarily constrains the real-world generalizability of the results. It is also possible that part of the observed performance advantage derives from the density of micro-adaptive updates within the meta-learning loop rather than from the quantum-inspired state representation itself. Future work should address these limitations through controlled real-learner studies, stronger empirical validation of biometric responsiveness, and more fine-grained ablation designs.

6.4. Broader Implications for Educational Technology

QCAL enters the domain of neurodata-driven education, where privacy, transparency, and autonomy are central. In the prototype, all physiological data are processed locally on edge devices, anonymized prior to any transmission, and governed by opt-in calibration with the option to disable tracking at any time, in line with GDPR and Türkiye's KVKK. To support human-in-the-loop oversight, QCAL exposes an analytics dashboard that visualizes cognitive-load and engagement traces alongside instructional adjustments (Rajpura et al., 2024; Guan et al., 2022). The framework thus represents not only a technical advance but also a human-centered approach that treats learners as active partners in personalization.

7. Conclusion

The Quantum-Cognitive Adaptive Learning (QCAL) framework brings together quantum-inspired cognition, quantum reinforcement learning, and biometric-informed adaptive feedback within a hybrid simulation-based architecture for personalized learning. Across the evaluated simulation scenarios and synthetic biometric signal conditions, QCAL outperformed the selected CALM baselines in learning speed, cognitive-load stabilization, and engagement recovery. The main contribution of the framework lies in representing learning as a probabilistic trajectory over a quantum-inspired state space (ψ), which can be dynamically updated through performance-related and biometric-style inputs under controlled computational conditions.

7.1. Policy and Practice Implications

Interdisciplinary collaboration among educational technology, cognitive science, learning analytics, and quantum computing researchers may help refine computationally informed adaptive learning systems. Privacy-sensitive design, transparent consent structures, and interpretable adaptation logic should remain central in any future attempt to translate biometric-informed personalization into real-world educational settings. Hybrid quantum-classical approaches may become more relevant as hardware and software ecosystems mature, but their educational value must be established through careful empirical validation rather than assumed from theoretical novelty alone.

7.2. Directions for Future Research

Future research should extend the present framework through controlled empirical studies with real learners, longitudinal validation, stronger sensor protocols, and cross-disciplinary benchmarking against advanced adaptive tutoring architectures. Additional work is also needed to clarify the relative contribution of quantum-inspired state modeling, biometric reward correction, and classical optimization components in realistic instructional environments.

8. Limitations and Future Work

Several limitations should be considered when interpreting the present findings. First, the evaluation is entirely simulation-based and therefore cannot fully capture the social, motivational, emotional, and contextual complexity of authentic educational environments. Second, the biometric component was validated using synthetic signal scenarios rather than data collected from real learners; accordingly, all conclusions regarding physiological adaptation remain preliminary. Third, although the hybrid quantum-classical design is conceptually grounded and computationally operationalized, the present results do not isolate with full certainty how much of the observed performance gain is attributable specifically to quantum-inspired modeling rather than to the broader adaptive control structure. An additional limitation concerns the use of a fixed deterministic seed configuration to maximize reproducibility and profile-level comparability between QCAL and CALM. Although this design facilitates controlled paired evaluation, future studies should incorporate multiple seed configurations and repeated simulation runs in order to quantify variability arising from initialization effects and stochastic adaptation trajectories.

Technical limitations should also be acknowledged. The proposed architecture assumes NISQ-compatible feasibility through shallow circuits and hybrid optimization, yet no real quantum hardware deployment was performed in the present study. In addition, several implementation choices, including ansatz structure, policy parameterization, and reward correction dynamics, remain open to further optimization and ablation under more realistic experimental settings.

Future validation should therefore proceed in stages: first through controlled laboratory studies with real learners and research-grade or carefully benchmarked consumer-grade sensors; second through longitudinal educational trials; and third through comparative testing against stronger adaptive learning baselines, including advanced LLM-supported tutoring and multimodal personalization systems. Such work will be necessary before stronger claims can be made regarding educational effectiveness, neuroadaptive validity, and deployable quantum-enhanced personalization.

8.1. Conflict of Interest Statement

The author declares that there is no conflict of interest regarding the publication of this article.

8.2. Funding Information

No external funding was received for conducting this study. The research was performed using the author's institutional resources.

8.3. Ethical Approval

This study relied entirely on simulated learner data and synthetic biometric signal scenarios. No human participants, human subjects, or animal subjects were involved, and no identifiable personal data were collected, processed, or stored. For that reason, formal ethics committee approval was not required for the present simulation-based design under standard institutional procedures. Any future extension involving real learners, live biometric sensing, or identifiable educational records would require prior ethics review and formal informed-consent procedures.

8.4. Data Availability Statement

The simulated learner profiles, synthetic biometric signal parameters, deterministic seed configuration (seed = 42), and analysis scripts underlying the present study can be made available by the corresponding author upon reasonable request. All data used in the study were artificially generated for simulation purposes and contain no personally identifiable information. The reported statistical contrasts were generated from paired synthetic learner profiles so that QCAL and CALM outcomes are comparable at the same learner-profile level.

References

Ajagekar, A., & You, F. (2024). Variational quantum circuit-based demand response in buildings leveraging a hybrid quantum-classical strategy. Applied Energy, 364, 123244. https://doi.org/10.1016/j.apenergy.2024.123244

Altinay, F., Karaatmaca, C., Altinay, Z., & Dagli, G. (2019). Developing at a great pace: Studies on artificial intelligence in higher education. eLearning and Software for Education, 330-337. https://doi.org/10.12753/2066-026X-19-116

Amato, F., Cicalese, M., Contrasto, L., Cubicciotti, G., D'Ambola, G., La Marca, A., ... Palomba, F. (2023). QuantuMoonLight: A low-code platform to experiment with quantum machine learning. SoftwareX, 22, 101399. https://doi.org/10.1016/j.softx.2023.101399

Balasubramani, R., & Biradar, V. G. (2024). Empowering autonomous decision-making through quantum reinforcement learning and cognitive neuromorphic frameworks. In 2024 4th International Conference on Mobile Networks and Wireless Communications (ICMNWC). IEEE. https://doi.org/10.1109/ICMNWC63764.2024.10872223

Baradari, D., Kosmyna, N., Petrov, O., Kaplun, R., & Maes, P.

(2025). NeuroChat: A neuroadaptive AI chatbot for customizing learning experiences. In Proceedings of the 7th ACM Conference on Conversational User Interfaces (CUI '25). Association for Computing Machinery. https://doi.org/10.1145/3719160.3736623

Bharti, K., Haug, T., Vedral, V., & Kwek, L.-C. (2020). Machine learning meets quantum foundations: A brief survey. AVS Quantum Science, 2(3), 034101. https://doi.org/10.1116/5.0007529

Buonaiuto, G., Guarasci, R., Minutolo, A., De Pietro, G., & Esposito, M. (2024). Quantum transfer learning for acceptability judgements. Quantum Machine Intelligence, 6(1), 13. https://doi.org/10.1007/s42484-024-00141-8

Chen, C., Zhu, Z., Zhou, X., & Chen, Q. (2008). Measurement-based quantum learning control of multi-qubit systems. Mediterranean Journal of Measurement and Control, 4(3), 94-100.

Chen, E., Lee, J.-E., Lin, J., & Koedinger, K. (2024). GPTutor: Great personalized tutor with large language models for personalized learning content generation. In Proceedings of the Eleventh ACM Conference on Learning@Scale (L@S '24) (pp. 539–541). Association for Computing Machinery. https://doi.org/10.1145/3657604.3664718

Cuellar, M. P., Pegalajar, M. C., & Cano, C. (2024). Automatic evolutionary design of quantum rule-based systems and applications to quantum reinforcement learning. Quantum Information Processing, 23(5), 179. https://doi.org/10.1007/s11128-024-04391-0

De Luca, G., & Chen, Y. (2023). Teaching quantum machine learning in computer science. In Proceedings of ISADS 2023 (pp. 189-195). IEEE. https://doi.org/10.1109/ISADS56919.2023.10092171

Dong, D., Chen, C., Li, H., & Tarn, T.-J. (2008). Quantum reinforcement learning. IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics), 38(5), 1207–1220. https://doi.org/10.1109/TSMCB.2008.925743

Dunjko, V., & Briegel, H. J. (2018). Machine learning and artificial intelligence in the quantum domain: A review of recent progress. Reports on Progress in Physics, 81(7), 074001. https://doi.org/10.1088/1361-6633/aab406

Esposito, M., Sabzevari, M. T., Ye, B., Falessi, D., Khan, A. A., & Taibi, D. (2024). Classi|Q>: Towards a translation framework to bridge the classical-quantum programming gap. In Proceedings of QSE-NE 2024 (pp. 11-14). ACM. https://doi.org/10.1145/3663531.3664752

Franklin, D., Palmer, J., Jang, W., Lehman, E. M., Marckwordt, J., Landsberg, R. H., Muller, A., & Harlow, D. (2020). Exploring quantum reversibility with young learners. In Proceedings of ICER 2020 (pp. 147-157). ACM. https://doi.org/10.1145/3372782.3406255

Ganesh, V., Seshia, S. A., & Jha, S. (2022). Machine learning and logic: A new frontier in artificial intelligence. Formal Methods in System Design, 60(3), 426-451. https://doi.org/10.1007/s10703-023-00430-1

Greengard, S. (2019). The algorithm that changed quantum machine learning. Communications of the ACM, 62(8), 15-17. https://doi.org/10.1145/3339458

Guan, J., Fang, W., & Ying, M. (2022). Verifying fairness in quantum machine learning. In Lecture Notes in Computer Science (Vol. 13372, pp. 408-429). Springer. https://doi.org/10.1007/978-3-031-13188-2_20

Hasanovic, M. (2023). Quantum education: How to teach a subject that nobody fully understands. In Proceedings of ETOP 2023 (Vol. 12723, 1272331). SPIE. https://doi.org/10.1117/12.2670468

Huang, R., Tan, X. Q., & Xu, Q. S. (2023). Learning to learn variational quantum algorithm. IEEE Transactions on Neural Networks and Learning Systems, 34(11), 8430-8440. https://doi.org/10.1109/TNNLS.2022.3151127

Jerbi, S., Gyurik, C., Marshall, S. C., Briegel, H. J., & Dunjko, V. (2021). Parametrized quantum policies for reinforcement learning. Advances in Neural Information Processing Systems, 34, 28362-28375. https://arxiv.org/abs/2103.05577

Koutromanos, D., Stefanatos, D., & Paspalakis, E. (2024). Control of qubit dynamics using reinforcement learning. Information, 15(5), 272. https://doi.org/10.3390/info15050272

Lamata, L. (2023). Quantum machine learning implementations: Proposals and experiments. Advanced Quantum Technologies, 6(7), 2300059. https://doi.org/10.1002/qute.202300059

Leporini, R. (2025). Extending a model language to handle entangled concepts in artificial intelligence. Foundations of Science, 30(3), 741–753. https://doi.org/10.1007/s10699-024-09956-x

Li, Z., Yazdanpanah, V., Wang, J., Gu, W., Shi, L., Cristea, A. I., Kiden, S., & Stein, S. (2025). TutorLLM: Customizing learning recommendations with knowledge tracing and retrieval-augmented generation. In Human-Computer Interaction - INTERACT 2025 (pp. 137-146). Springer. https://doi.org/10.1007/978-3-032-05005-2_8

Maksimovic, M., & Maksymov, I. S. (2025). Transforming neural networks into quantum-cognitive models: A research tutorial with novel applications. Technologies, 13(5), 183. https://doi.org/10.3390/technologies13050183

Martin-Guerrero, J. D., & Lamata, L. (2022). Quantum machine learning: A tutorial. Neurocomputing, 470, 457-461. https://doi.org/10.1016/j.neucom.2021.02.102

Moreira, C., Fell, L., Dehdashti, S., Bruza, P., & Wichert, A. (2020). Towards a quantum-like cognitive architecture for decision-making. Behavioral and Brain Sciences, 43, e17. https://doi.org/10.1017/S0140525X19001687

Neumann, N. M. P., de Heer, P. B. U. L., & Phillipson, F. (2023). Quantum reinforcement learning: Comparing quantum annealing and gate-based quantum computing with classical deep reinforcement learning. Quantum Information Processing, 22(2), 125. https://doi.org/10.1007/s11128-023-03867-9

Pothos, E. M., & Busemeyer, J. R. (2022). Quantum cognition. Annual Review of Psychology, 73, 749-778. https://doi.org/10.1146/annurev-psych-033020-123501

Rajpura, P., Cecotti, H., & Meena, Y. K. (2024). Explainable artificial intelligence approaches for brain–computer interfaces: A review and design space. Journal of Neural Engineering, 21(4), 041003. https://doi.org/10.1088/1741-2552/ad6593

Reinikainen, J., Stirbu, V., Heinosaari, T., Lappalainen, V., & Mikkonen, T. (2024). Quantum computing for all: Online courses built around an interactive visual quantum circuit simulator. IEEE Computer Graphics and Applications, 44(5), 67-75. https://doi.org/10.1109/MCG.2024.3444952

Ruan, S., Liang, Z., Guan, Q., Griffin, P., Wen, X., Lin, Y., & Wang, Y. (2024). VIOLET: Visual analytics for explainable quantum neural networks. IEEE Transactions on Visualization and Computer Graphics, 30(6), 2862-2874. https://doi.org/10.1109/TVCG.2024.3388557

Salehi, O., Seskir, Z., & Tepe, I. (2022). A computer-science-oriented approach to introduce quantum computing to a new audience. IEEE Transactions on Education, 65(1), 1-8. https://doi.org/10.1109/TE.2021.3078552

Sequeira, A., Santos, L. P., & Barbosa, L. S. (2024). On quantum natural policy gradients. IEEE Transactions on Quantum Engineering, 5, 3102411. https://doi.org/10.1109/TQE.2024.3418094

Shahwar, T., Zafar, J., Almogren, A., Zafar, H., Rehman, A. U., Shafiq, M., & Hamam, H. (2022). Automated detection of Alzheimer's via hybrid classical-quantum neural networks. Electronics, 11(5), 721. https://doi.org/10.3390/electronics11050721

Signorelli, C. M., & Arsiwalla, X. D. (2019). Moral dilemmas for artificial intelligence: A position paper on an application of compositional quantum cognition. In Lecture Notes in Computer Science (Vol. 11690, pp. 123-138). Springer. https://doi.org/10.1007/978-3-030-35895-2_9

Stokes, J., Izaac, J., Killoran, N., & Carleo, G. (2020). Quantum natural gradient. Quantum, 4, 269. https://doi.org/10.22331/q-2020-05-25-269

Vedaie, S. S., Dalal, A., Paez, E. J., & Sanders, B. C. (2023). Framework for learning and control in the classical and quantum domains. Annals of Physics, 458, 169471. https://doi.org/10.1016/j.aop.2023.169471

Yan, F., Iliyasu, A. M., Li, N., Salama, A. S., & Hirota, K. (2024). Quantum robotics: A review of emerging trends. Quantum Machine Intelligence, 6(2), 86. https://doi.org/10.1007/s42484-024-00225-5

Zuccarini, G., Malgieri, M., Sutrini, C., & Macchiavello, C. (2024). Elementarizing quantum algorithms: Clarification of the internal structure and preliminary learning outcomes. Journal of Physics: Conference Series, 2750, 012025. https://doi.org/10.1088/1742-6596/2750/1/012025