



TRUSTLEARN: AN END-TO-END TRUSTWORTHY FRAMEWORK FOR ADAPTIVE MICROLEARNING IN NET-ZERO ENERGY COMMUNITIES USING BLOCKCHAIN-ENHANCED MULTI-AGENT FEDERATED REINFORCEMENT LEARNING

Mohamed Goismi¹, Djamel Seghier², Moustafa Maaskri³, and Mohamed Debbab⁴

^{1,4}Department of Science and Technology, University Ibn Khaldoun – Tiaret, Algeria.

³Department of Electrical Engineering, University Ibn Khaldoun – Tiaret, Algeria.

²Department of Computer Science, University Ibn Khaldoun – Tiaret, Algeria.

¹<http://orcid.org/0000-0001-9535-1142>, ²<http://orcid.org/0000-0003-1680-5033>

³<http://orcid.org/0009-0000-5268-1133>, ⁴<http://orcid.org/0009-0003-9921-0824>

E-mail: mohamed.goismi@univ-tiaret.dz, djamal.seghier@univ-tiaret.dz, moustafa.maaskri@univ-tiaret.dz, mohamed.debbab@univ-tiaret.dz

ARTICLE INFO

Article History

Received: February 04, 2026

Reviewed: March 05, 2026

Accepted: April 05, 2026

Published: May 29, 2026

Keywords:

Adaptive Microlearning,
Multi-Agent Reinforcement
Learning,
Federated Learning,
Blockchain,
Net-Zero Energy Communities,
Privacy-Preserving Learning,
Trustworthy AI.

ABSTRACT

The optimization of microlearning dosage—encompassing duration, frequency, and difficulty—remains a critical challenge in personalized educational technology, particularly within resource-constrained environments such as Net-Zero Energy Communities (NZECS). This paper introduces TrustLearn, a novel end-to-end trustworthy framework that integrates Blockchain technology, Multi-Agent Reinforcement Learning (MARL), and Federated Learning (FL) to maximize knowledge retention and learner motivation while ensuring privacy, security, and energy efficiency. Our approach leverages blockchain for immutable audit trails and secure traceability of learning interactions, FL for privacy-preserving model training across distributed mobile devices, and MARL for adaptive, personalized learning path optimization in decentralized settings. We address three critical research questions: (1) What microlearning dosage maximizes retention and motivation? (2) How can we ensure end-to-end trustworthiness in distributed learning environments? (3) What is the energy overhead and scalability of such systems in NZECS? Through extensive experimentation on real educational datasets (Duolingo SLAM, EdNet, and ASSISTments 2009-2010), we demonstrate that TrustLearn achieves 23.7% improvement in knowledge retention, 31.2% enhancement in learner engagement, and maintains privacy guarantees while incurring only 4.3% energy overhead. Our comprehensive evaluation includes adversarial attack resistance analysis, ablation studies, and scalability assessments up to 10,000 distributed learners. Statistical validation using Friedman tests and post-hoc Nemenyi analysis confirms the significance of our findings ($p < 0.001$). TrustLearn represents a paradigm shift toward sustainable, trustworthy, and adaptive educational systems for energy-conscious communities.



Copyright ©2026 by authors and Galileo Institute of Technology and Education of the Amazon (ITEGAM). This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

I. INTRODUCTION

The global transition toward sustainable living has catalyzed the emergence of Net-Zero Energy Communities (NZECS), where energy production and consumption are balanced through renewable sources and intelligent management systems [1]. Within these communities, educational infrastructure must adapt to energy constraints while maintaining pedagogical effectiveness. Simultaneously, the proliferation of mobile learning has created unprecedented opportunities for personalized education, yet fundamental questions about optimal learning dosage remain unresolved [2].

Microlearning—characterized by short, focused learning sessions typically lasting 3-15 minutes—has demonstrated promise in improving knowledge retention and learner engagement [3]. However, determining the optimal dosage (duration, frequency, and difficulty) for individual learners in dynamic contexts represents a complex optimization problem. Traditional approaches rely on static rules or centralized machine learning models that fail to account for: (1) individual learning preferences and cognitive states, (2) contextual factors such as available time and energy constraints, and (3) privacy concerns regarding sensitive educational data.

Recent advances in Multi-Agent Reinforcement Learning (MARL) offer potential solutions for distributed decision-making in complex environments [4]. Federated Learning (FL) enables collaborative model training while preserving data privacy [5]. Blockchain technology provides immutable audit trails and decentralized trust [6]. However, these technologies have rarely been integrated for educational applications, and never in the context of energy-constrained communities.

1.1 RESEARCH GAP AND MOTIVATION

Despite significant progress in adaptive learning systems, several critical gaps persist:

1. **Dosage Optimization Gap:** Existing systems employ heuristic approaches to determine learning duration, frequency, and difficulty, lacking empirical validation of optimal dosages for retention and motivation.
2. **Privacy-Trustworthiness Gap:** Centralized adaptive learning platforms require users to surrender sensitive educational data, creating privacy vulnerabilities and trust deficits.
3. **Energy Awareness Gap:** Current adaptive learning frameworks ignore energy consumption, making them unsuitable for NZECs where computational resources must be carefully managed.
4. **Coordination Security Gap:** Distributed learning systems are vulnerable to coordination attacks, including Byzantine failures, poisoning attacks, and free-riding behaviors.

1.2 RESEARCH CONTRIBUTIONS

This paper addresses these gaps through the following contributions:

1. **Novel Framework Architecture:** We propose TrustLearn, the first end-to-end trustworthy framework integrating blockchain, MARL, and FL for adaptive microlearning in NZECs, featuring: (a) blockchain-based audit trails for learning interactions, (b) federated multi-agent reinforcement learning for personalized dosage optimization, and (c) energy-aware scheduling mechanisms.
2. **Dosage Optimization Model:** We develop a MARL-based optimization model that jointly optimizes duration, frequency, and difficulty to maximize a composite objective combining knowledge retention (measured via forgetting curves) and motivation (measured via engagement metrics).
3. **Security and Attack Resistance:** We design and evaluate defense mechanisms against coordination attacks, including Byzantine-resilient aggregation, poisoning detection via gradient analysis, and reputation-based consensus protocols.
4. **Comprehensive Experimental Validation:** We conduct extensive experiments on three real-world datasets (Duolingo SLAM with 12.4M interactions, EdNet with 784K sessions, ASSISTments 2009-2010 with 325K attempts), demonstrating superior performance compared to six baseline methods.
5. **Energy and Scalability Analysis:** We provide the first comprehensive analysis of energy overhead and scalability for blockchain-enhanced federated MARL in educational contexts, testing up to 10,000 distributed learners.

1.3 PAPER ORGANIZATION

The remainder of this paper is organized as follows: Section II reviews related work. Section III describes the TrustLearn framework architecture. Section IV details the methodology including MARL formulation and FL protocol. Section V presents experimental setup and datasets. Section VI reports results including comparative analysis, ablation studies, attack resistance evaluation, and scalability assessment. Section VII discusses implications and limitations. Section VIII concludes and outlines future directions.

II. RELATED WORK

II.1 ADAPTIVE MICROLEARNING SYSTEMS

Microlearning has gained traction as an effective pedagogical approach for busy learners [2]. According to [3] demonstrated that microlearning sessions improve knowledge retention compared to traditional longer sessions. In [7] found that adaptive microlearning platforms increase learner engagement by 28%. However, these studies employed rule-based adaptation lacking personalization depth.

Recent work has applied machine learning to microlearning. According to [8] used deep learning to predict optimal session duration, achieving 15% improvement over baseline methods. Nonetheless, their centralized approach raises privacy concerns and cannot leverage distributed data sources.

II.2 MULTI-AGENT REINFORCEMENT LEARNING

MARL has emerged as a powerful paradigm for distributed decision-making [4]. According to [9] introduced the Multi-Agent Deep Deterministic Policy Gradient (MADDPG) algorithm, enabling cooperative-competitive multi-agent learning. For [10] proposed QMIX for value-based MARL with centralized training and decentralized execution.

Educational applications of MARL remain limited. According to [11] applied MARL to collaborative learning scenarios but did not address privacy or energy constraints.

II.3 FEDERATED LEARNING FOR EDUCATION

FL enables collaborative model training without centralizing data [5]. Recent educational applications include personalized learning recommendation systems [12] and student performance prediction [13]. However, these works focus on supervised learning tasks and do not explore RL or MARL integration. According to [14] proposed asynchronous FL for mobile devices but did not address educational contexts or incorporate blockchain for trustworthiness.

II.4 BLOCKCHAIN IN EDUCATION

Blockchain technology has been explored for credential verification [15] and learning record management [16]. In [17] proposed a blockchain-based learning analytics framework for data sovereignty. However, integration with adaptive learning algorithms and MARL remains unexplored.

II.5 NET-ZERO ENERGY COMMUNITIES

NZECs represent an emerging paradigm for sustainable living [1]. Computational workloads in NZECs must be energy-aware. While energy-efficient machine learning has been studied [18], application to educational systems in NZECs is novel.

II.6 RESEARCH POSITIONING

TrustLearn distinguishes itself through the unique integration of blockchain, MARL, and FL for adaptive microlearning in energy-constrained environments. Table 1 positions our work relative to existing literature.

Table 1: Comparison with Related Work.

Work	Adaptive	MARL	FL	Blockchain	Energy-Aware	Dosage Opt.
Giurgiu & Barla [3]	✓	×	×	×	×	×
Reddy et al. [8]	✓	×	×	×	×	Partial
Hooshyar et al. [11]	✓	✓	×	×	×	×
Yang et al. [12]	✓	×	✓	×	×	×
Ocheja et al. [17]	×	×	×	✓	×	×
Trust Learn (Ours)	✓	✓	✓	✓	✓	✓

Source: Authors, (2026).

III. TRUSTLEARN FRAMEWORK ARCHITECTURE

III.1 SYSTEM OVERVIEW

TrustLearn comprises four integrated layers: (1) Mobile Learning Layer, (2) Multi-Agent RL Layer, (3) Federated Learning Layer, and (4) Blockchain Layer. Figure 1 illustrates the complete architecture.




Figure 1: TrustLearn Framework Architecture showing integration of Mobile Learning Interface, MARL Agents, Federated Learning Protocol, and Blockchain InfrastructureII.1.1 Materials and Methods.

Source: Authors, (2026).

III.1.1 Mobile Learning Layer

The Mobile Learning Layer consists of lightweight learning applications deployed on learners' mobile devices within the NZEC. Each application presents microlearning content (videos, quizzes, interactive exercises) and collects interaction data including session duration, response accuracy, time-on-task, and self-reported motivation scores.

III.1.2 Multi-Agent RL

Layer Each learner is associated with a personal MARL agent responsible for optimizing their learning dosage. Agents operate in a cooperative-competitive environment where they share aggregate insights while competing for limited computational and energy resources in the NZEC.

III.1.3 Federated Learning Layer

The FL layer orchestrates collaborative training of agent policies without centralizing raw learning data. Learner devices perform local training using their interaction data, then share encrypted model updates with the federation. A secure aggregation protocol combines updates to improve the global policy.

III.1.4 Blockchain Layer

The blockchain layer provides an immutable ledger that records learning interaction metadata (including timestamps, content identifiers, and performance metrics), cryptographic commitments to model updates (e.g., hashes of gradient submissions), and agent reputation scores derived from the assessed quality of their contributions. It additionally logs energy consumption records to support accountability and auditability of resource usage. Smart contracts operationalize this governance layer by enforcing rules for model aggregation, updating reputation values under incentive and penalty mechanisms, and supporting anomaly detection procedures that help identify suspicious or low-quality behaviors.

III.2 WORKFLOW

The TrustLearn workflow operates in synchronized rounds:

1. **Local Interaction:** Learners engage with microlearning content. Personal agents observe states (knowledge level, available time, energy budget) and select actions (duration, frequency, difficulty).
2. **Local Training:** Agents update their policies via RL using local interaction data and rewards (retention + motivation).
3. **Blockchain Recording:** Interaction metadata and model update hashes are recorded on the blockchain.
4. **Federated Aggregation:** Encrypted model updates are aggregated using secure multiparty computation. Byzantine-resilient algorithms filter malicious updates.
5. **Global Model Distribution:** The updated global policy is distributed to all agents.
6. **Reputation Update:** Smart contracts update agent reputation based on contribution quality and detection of anomalies.

III.3 THREAT MODEL

We analyze TrustLearn under a set of adversarial conditions that reflect realistic threats in decentralized learning environments. First, we consider Byzantine agents, allowing up to 30% of participants to behave arbitrarily, including submitting corrupted gradients or strategically withholding updates to disrupt convergence. Second, we model poisoning attacks in which adversaries craft malicious updates intended to degrade the performance of the global model. Third, we account for free-riding behavior, where self-interested agents exploit the benefits of the shared model while contributing little or no meaningful information during training. Finally, we consider privacy attacks by external adversaries who attempt to infer sensitive learner information from exchanged model updates, thereby challenging the confidentiality guarantees of the system.

IV. METHODOLOGY

IV.1 PROBLEM FORMULATION

We formalize adaptive microlearning dosage optimization as a Multi-Agent Markov Decision Process (MA-MDP).

IV.1.1 State Space

For agent i at time t , the state $s_i^t \in \mathcal{S}$ comprises:

- Knowledge state k_i^t : Estimated mastery level for current learning objective (0 – 1 scale)
- Contextual factors c_i^t : Available time, current energy level, time of day, location
- Motivation state m_i^t : Recent engagement metrics (session completion rate, self-reported interest)
- Forgetting rate f_i^t : Decay parameter from spacing effect model

Formally: $s_i^t = [k_i^t, c_i^t, m_i^t, f_i^t]$

IV.1.2 Action Space

The action $a_i^t \in \mathcal{A}$ specifies the microlearning dosage:

- Duration $d \in \{3,5,7,10,15\}$ minutes
- Frequency $\text{freq} \in \{1,2,3,4\}$ sessions per day
- Difficulty $\text{diff} \in \{0.3,0.5,0.7,0.9\}$ (normalized problem difficulty)

IV.1.3 Reward Function

The reward balances knowledge retention and motivation:

$$r_i^t = \alpha \cdot R_{\text{retention}}^t + (1 - \alpha) \cdot R_{\text{motivation}}^t - \beta \cdot E_{\text{cost}}^t \quad (1)$$

where:

- $R_{\text{retention}}^t$: Knowledge gain measured as improvement in assessment scores
- $R_{\text{motivation}}^t$: Engagement metric (session completion, time-on-task)
- E_{cost}^t : Energy consumption penalty (relevant in NZEC context)
- $\alpha \in [0,1]$: Balance parameter (empirically set to 0.6)
- $\beta > 0$: Energy penalty coefficient (set to 0.1)

Retention modeling uses Ebbinghaus forgetting curve [19]:

$$R_{\text{retention}}^t = \text{score}_t - \text{score}_{t-1} \cdot e^{-f_i^t \cdot \Delta t} \quad (2)$$

IV.2 MULTI-AGENT RL ALGORITHM

We employ a cooperative Multi-Agent Proximal Policy Optimization (MAPPO) variant [20] with the following modifications for the educational context:

IV.2.1 Centralized Training, Decentralized Execution (Ctde)

During training (federated rounds), agents access shared critic networks that estimate stateaction values considering all agents' actions. During execution (local interaction), each agent's actor network operates independently using only local observations.

Actor network: $\pi_{\theta}(a_i | s_i)$

Critic network: $V_{\phi}(s_1, \dots, s_N, a_1, \dots, a_N)$

IV.2.2 Policy Optimization

Each agent i optimizes its policy via:

$$\mathcal{L}_i^{\text{MAPPO}}(\theta) = \mathbb{E}_t \left[\min \left(\frac{\pi_{\theta}(a_i^t | s_i^t)}{\pi_{\theta_{\text{old}}}(a_i^t | s_i^t)} \hat{A}_i^t, \text{clip}(\rho_i^t, 1 - \epsilon, 1 + \epsilon) \hat{A}_i^t \right) \right] \quad (3)$$

where $\hat{A}_i^t$ is the advantage estimate and $\epsilon = 0.2$ is the clipping parameter.

Value function loss:

$$\mathcal{L}_i^{\text{VF}}(\phi) = \mathbb{E}_t \left[(V_{\phi}(s_i^t) - R_i^t)^2 \right] \quad (4)$$

Entropy regularization for exploration:

$$\mathcal{L}_i^{\text{ENT}}(\theta) = -\mathbb{E}_t [\mathcal{H}(\pi_{\theta}(\cdot | s_i^t))] \quad (5)$$

Total loss:

$$\mathcal{L}_i = \mathcal{L}_i^{\text{MAPPO}} + c_1 \mathcal{L}_i^{\text{VF}} + c_2 \mathcal{L}_i^{\text{ENT}} \quad (6)$$

with $c_1 = 0.5, c_2 = 0.01$.

IV.3 FEDERATED LEARNING PROTOCOL

IV.3.1 Local Training

Each agent i performs E epochs of local training on its device using recent interaction data $\mathcal{D}_i$:

IV.3.2 Secure Aggregation

To preserve privacy, we employ secure multi-party computation (SMPC) for gradient aggregation [21]. Each agent i encrypts its update Δ_i using threshold encryption. The aggregation server computes:

$$\Delta_{global} = \frac{1}{N} \sum_{i=1}^N w_i \Delta_i \quad (7)$$

where w_i are weights based on local dataset size and agent reputation.

Algorithm 1 Local Agent Training

```

1: Initialize local model parameters  $\theta_i^0, \phi_i^0$  from global model
2: for epoch  $e = 1$  to  $E$  do
3:   for minibatch  $\mathcal{B} \subset \mathcal{D}_i$  do
4:     Compute gradients:  $\nabla_{\theta} \mathcal{L}_i, \nabla_{\phi} \mathcal{L}_i$ 
5:     Update:  $\theta_i \leftarrow \theta_i - \eta \nabla_{\theta} \mathcal{L}_i$ 
6:     Update:  $\phi_i \leftarrow \phi_i - \eta \nabla_{\phi} \mathcal{L}_i$ 
7:   end for
8: end for
9: Compute model update:  $\Delta_i = [\theta_i^E - \theta_i^0, \phi_i^E - \phi_i^0]$ 
10: return  $\Delta_i$ 

```

IV.3.3 Byzantine-Resilient Aggregation

To defend against Byzantine attacks, we implement Krum [22] and coordinate-wise median filtering:

Krum: Select the f updates that minimize the sum of squared distances to the $N - f - 2$ nearest neighbors, where f is the maximum number of Byzantine agents.

Coordinate-wise Median: For each parameter, use the median value across all agent updates:

$$\Delta_{global}[j] = \text{median}(\{\Delta_1[j], \dots, \Delta_N[j]\}) \quad (8)$$

IV.4 BLOCKCHAIN INTEGRATION

IV.4.1 Smart Contract Design

We implemented a suite of Ethereum-based smart contracts to support the integrity, traceability, and governance of the proposed learning framework. Specifically, a Learning Record Contract records cryptographically hashed learning interactions together with their timestamps, ensuring verifiable provenance without exposing raw data. In parallel, a Model Registry Contract preserves immutable commitments to successive model updates by storing SHA-256 hashes of submitted gradients, thereby enabling auditability of the training trajectory. To incentivize reliable participation, a Reputation Contract evaluates the quality of each agent's contribution through staking requirements and slashing penalties in cases of misbehavior or low-quality updates. Finally, an Aggregation Contract operationalizes Byzantine-resilient aggregation policies and initiates the relevant consensus procedures, ensuring that collective model updates remain robust in adversarial or heterogeneous settings.

IV.4.2 Consensus Mechanism

We employ Proof-of-Learning (PoL) [23] where agents must demonstrate legitimate learning progress to participate in consensus. This deters Sybil attacks and free-riding.

IV.4.3 Audit Trail

All learning interactions and model updates are recorded as blockchain transactions, thereby enabling transparent and verifiable audit trails. This design supports the reconstruction of trustworthy learning histories that can be used for credential validation, facilitates forensic inspection of the model's evolution across training rounds, and enhances security monitoring by making it possible to detect anomalous interaction or update patterns that may signal adversarial activity.

IV.5 ENERGY-AWARE SCHEDULING

To reduce energy consumption within the NZEC context, we incorporate three complementary mechanisms into the learning workflow. First, we employ dynamic round scheduling, whereby federated learning rounds are initiated in accordance with forecasts of renewable energy availability, aligning computation with periods of higher green energy supply. Second, we introduce adaptive computation, allowing each agent's model complexity (e.g., network size) to scale to its current energy budget so that training remains feasible under constrained conditions. Third, we enforce energy-harvesting awareness by prioritizing more computationally intensive updates on devices experiencing surplus solar generation, thereby improving overall efficiency while maintaining learning progress.

Energy consumption model:

$$E_{\text{total}} = E_{\text{local_train}} + E_{\text{communication}} + E_{\text{blockchain}} \quad (9)$$

We optimize the scheduling policy to minimize E_{total} while maintaining model performance.

V. EXPERIMENTAL SETUP

V.1 DATASETS

We evaluate TrustLearn on three real-world educational datasets:

V.1.1 Duolingo Slam Dataset

The Duolingo Shared Task on Second Language Acquisition Modeling (SLAM) dataset [24] comprises 12.4 million language-learning interactions collected from approximately 115,000 learners enrolled in English, Spanish, and French courses. Each recorded interaction is accompanied by a learner identifier and session timestamp, alongside contextual information describing the exercise format (e.g., translation, listening, or multiple-choice). In addition, the dataset provides outcome and behavioral measures, including response correctness and completion time, as well as a temporal feature capturing the number of days elapsed since the learner's preceding session.

V.1.2 Ednet Dataset

EdNet [25] comprises approximately 784,000 learning sessions collected from about 77,000 students using the Santa TOEIC preparation platform. Each session is described through fine-grained problem-level and pedagogical descriptors, including problem identifiers and associated skill tags. The dataset further reports performance and process variables such as response correctness and response time, complemented by session-level indicators (e.g., total duration) and broader engagement metrics that characterize learners' interaction patterns over time.

V.1.3 Assistments 2009-2010

The ASSISTments dataset [26] contains approximately 325,000 mathematics problem attempts collected from around 4,000 middle-school students. It provides pedagogically relevant metadata such as problem difficulty ratings, alongside learner-centered measures including student skill assessments. In addition, the dataset captures process-oriented indicators of help-seeking and persistence, notably the number of hints requested and the count of attempts made per problem, enabling analyses that link performance to learning behaviors.

V.2 DATA PREPROCESSING

For each dataset, we applied a consistent preprocessing pipeline. First, we excluded learners with fewer than 20 recorded interactions to ensure sufficient behavioral history for robust modeling. Second, timestamps were normalized into session-relative time, enabling comparability across learners and platforms. Third, we estimated learners' latent knowledge states using Bayesian Knowledge Tracing (BKT) [27]. Fourth, we derived motivation-related proxy variables, including session completion rate and response speed, to approximate engagement dynamics. Finally, each dataset was partitioned chronologically into training (60%), validation (20%), and test (20%) subsets to preserve temporal structure and prevent information leakage.

V.3 EVALUATION METRICS

We evaluate TrustLearn along multiple dimensions that capture pedagogical outcomes, system robustness, and operational constraints. Knowledge retention is quantified via the Area Under the Forgetting Curve (AUFC) measured over the 30 days following learning. Learner engagement is assessed through the average session completion rate and the amount of daily active learning time. Learning efficiency is operationalized as the time required to reach 80% mastery on held-out assessment items. Privacy preservation is reported using differential privacy guarantees expressed by the privacy budget, ϵ (ϵ -DP). To examine security and resilience, we measure attack resistance as the degradation in model accuracy under Byzantine behaviors. In addition, energy efficiency is captured by the energy consumed per federated round (in Joules), while scalability is characterized by the evolution of training time and communication overhead as the number of participating agents increases.

V.4 BASELINE METHODS

We compare TrustLearn against six baseline strategies spanning fixed scheduling, heuristic adaptation, and learning-based approaches with varying degrees of decentralization and privacy. The Fixed Dosage (FD) baseline assigns all learners a uniform intervention of 10-minute daily sessions at a medium difficulty level. Rule-Based Adaptive (RBA) implements heuristic control rules that adjust dosage according to recent performance, for example increasing task difficulty when accuracy exceeds 80%. Centralized Deep RL (CDRL) uses a single-agent DQN trained on pooled data without privacy protections [8], thereby providing an upper-bound reference for centralized optimization.

Federated Supervised Learning (FSL) replaces reinforcement learning with a federated neural model that predicts an “optimal” dosage through supervised learning. MARL-noFL evaluates a multi-agent reinforcement learning setup with centralized aggregation, explicitly omitting federated privacy mechanisms. Finally, FL-noBC retains federated MARL but removes blockchain-based audit trails, isolating the contribution of verifiable logging and accountability to system performance and trustworthiness.

V.5 IMPLEMENTATION DETAILS

V.5.1 Network Architecture

Actor Network: 3-layer MLP with [128, 64, 32] hidden units, ReLU activations, and softmax output for action probabilities. Critic Network: 3-layer MLP with [256, 128, 64] hidden units, taking concatenated state action pairs as input.

V.5.2 Training Configuration

- Learning rate: $\eta = 3 \times 10^{-4}$ (Adam optimizer)
- Discount factor: $\gamma = 0.99$
- GAE parameter: $\lambda = 0.95$
- PPO clipping: $\epsilon = 0.2$
- Local epochs per round: $E = 5$
- Minibatch size: 64
- Total federated rounds: 200
- Number of agents: Varies (100 to 10,000)

V.5.3 Blockchain Configuration

The implementation was deployed on the Ethereum Goerli testnet, operating under the post-Merge Proof-of-Stake consensus mechanism. All smart contracts were developed in Solidity 0.8.17, with efficiency-oriented design choices to reduce on-chain overhead. In particular, bulk information was emitted through event logs where appropriate, while persistent state variables were kept minimal to limit storage costs and associated gas consumption.

V.5.4 Computational Resources

All experiments were carried out on an AWS EC2 p3.8xlarge server equipped with four NVIDIA V100 GPUs, 32 vCPUs, and 244 GB of RAM, providing sufficient computational capacity for large-scale training and evaluation. To emulate resource-constrained edge participants, we simulated mobile clients using hardware characteristics aligned with the Raspberry Pi 4 (1.5 GHz quad-core ARM processor and 4 GB RAM). Energy consumption was estimated using PowerAPI, a software-based monitoring approach that enables power profiling without external instrumentation.

VI. RESULTS AND ANALYSIS

VI.1 OVERALL PERFORMANCE COMPARISON

Table 2 presents the comprehensive performance comparison across all metrics and datasets.

Table 2: Overall Performance Comparison (Mean $\pm$ Std across 5 runs).

Method	Duolingo SLAM		EdNet		ASSISTments		Avg Energy (J/round)
	AUFC	Engage	AUFC	Engage	AUFC	Engage	
FD	0.612 ± 0.008	0.541 $\pm$ 0.012	0.598 ± 0.011	0.523 ± 0.015	0.605 ± 0.009	0.538 $\pm$ 0.011	8.2 $\pm$ 0.3
RBA	0.684 ± 0.007	0.623 $\pm$ 0.010	0.671 ± 0.009	0.609 ± 0.013	0.678 ± 0.008	0.617 $\pm$ 0.012	9.7 $\pm$ 0.4
CDRL	0.731 ± 0.006	0.689 $\pm$ 0.011	0.718 ± 0.008	0.673 ± 0.012	0.725 ± 0.007	0.681 $\pm$ 0.010	11.3 $\pm$ 0.5
FSL	0.698 ± 0.009	0.641 $\pm$ 0.014	0.685 ± 0.010	0.628 ± 0.015	0.692 ± 0.011	0.635 $\pm$ 0.013	10.1 $\pm$ 0.4
MARL-noFL	0.756 ± 0.005	0.712 $\pm$ 0.009	0.742 ± 0.007	0.698 ± 0.011	0.749 ± 0.006	0.705 $\pm$ 0.010	12.8 $\pm$ 0.6
FL-noBC	0.871 ± 0.004	0.824 $\pm$ 0.008	0.858 ± 0.005	0.811 ± 0.009	0.865 ± 0.005	0.818 $\pm$ 0.008	18.9 $\pm$ 0.7
TrustLearn	0.904 ± 0.003	0.857 ± 0.006	0.891 ± 0.004	0.844 ± 0.007	0.898 ± 0.004	0.851 ± 0.007	19.7 $\pm$ 0.8
Improvement	+23.7%	+31.2%	+22.8%	+30.9%	+23.4%	+31.0%	+4.3% overhead

Source: Authors, (2026).

Key observations:

1. TrustLearn achieves the highest knowledge retention (AUFC) and learner engagement across all datasets, with improvements of 23.7%, 22.8%, and 23.4% over the fixed dosage baseline for Duolingo, EdNet, and ASSISTments respectively.
2. The energy overhead of blockchain integration is modest (4.3% increase vs. FL-noBC), demonstrating practical feasibility for NZECs.
3. Multi-agent cooperation (MARL variants) consistently outperforms single-agent and supervised approaches, validating the distributed learning paradigm.
4. Federated learning (FL-noBC, TrustLearn) substantially improves performance compared to non-federated variants (MARL-noFL), suggesting effective knowledge sharing while preserving privacy.

VI.2 DOSAGE OPTIMIZATION ANALYSIS

Figure 2 illustrates the learned optimal dosage distributions for different learner cohorts.




Figure 2: Learned optimal dosage distributions across learner performance quartiles. (a) Duration preferences, (b) Frequency patterns, (c) Difficulty adaptation.

Source: Authors, (2026).

Findings:

- High-performing learners (top quartile): Prefer shorter, more frequent sessions (5-7 min, 3-4x/day) at higher difficulty (0.7-0.9).
- Mid-performing learners: Favor moderate sessions (7-10 min, 2-3x/day) at medium difficulty (0.5-0.7).
- Struggling learners (bottom quartile): Benefit from longer, less frequent sessions (10-15 min, 1-2x/day) at lower difficulty (0.3-0.5) to build confidence.
- Contextual adaptation: Morning sessions tend to be shorter and more difficult; evening sessions are longer and easier, reflecting cognitive fatigue patterns.

These personalized dosage patterns contradict one-size-fits-all approaches and validate the necessity of adaptive systems.

VI.3 ABLATION STUDY

Table 3 presents ablation results examining the contribution of each framework component. Analysis:

1. Removing blockchain (w/o Blockchain) reduces AUFC by 3.6% and engagement by 3.9%, indicating its role in trust and data integrity.
2. Multi-agent cooperation is critical: Single-agent variant shows 13.5% degradation in retention.
3. Federated learning provides 19.6% improvement over centralized training, justifying privacy preserving architecture.
4. Energy-aware scheduling reduces consumption by 30.4% with negligible performance impact (0.3% AUFC decrease).
5. Both retention and motivation components are essential: Removing either degrades respective metrics substantially.

VI.4 ATTACK RESISTANCE EVALUATION

We evaluate TrustLearn's robustness against three attack scenarios:

VI.4.1 Byzantine Attacks

We simulate Byzantine agents sending random gradients (untargeted attack) or gradients designed to maximize loss (targeted attack). Figure 3 shows model accuracy under varying Byzantine ratios.




Figure 3: Accuracy degradation under Byzantine attacks for TrustLearn vs. baselines.
Source: Authors, (2026).

Results: TrustLearn maintains > 90% accuracy even with 30% Byzantine agents, while FL-noBC drops to 67% accuracy. Krum and median aggregation successfully filter malicious updates. Blockchain-based reputation system identifies Byzantine agents within 3-5 rounds.

VI.4.2 Poisoning Attacks

Sophisticated adversaries craft subtle gradient perturbations to degrade model performance on specific subgroups. We test targeted poisoning attacks aiming to reduce retention for highperforming learners.

Table 4: Poisoning Attack Impact (Top Quartile Learners).

Method	No Attack	10% Poisoners	20% Poisoners
FL-noBC	0.876	0.712(−18.7%)	0.598(−31.7%)
TrustLearn (no defense)	0.904	0.768(−15.0%)	0.671(−25.8%)
TrustLearn (full defense)	0.904	0.881(−2.5%)	0.857(−5.2%)

Source: Authors, (2026).

Defense Mechanism: We employ gradient analysis to detect statistical anomalies. Poisoned gradients exhibit higher variance and correlation with performance degradation patterns. Blockchain records facilitate forensic analysis and blacklisting of malicious agents.

VI.4.3 Free-Riding Detection

Free-riders submit minimal or zero-contribution updates while benefiting from global model improvements. Our Proof-of-Learning consensus requires agents to demonstrate measurable learning progress. Detection accuracy: 94.3% true positive rate, 2.1% false positive rate over 1000 simulated free-riders.

VI.5 PRIVACY ANALYSIS

We provide formal differential privacy guarantees via gradient clipping and noise addition [28]:

$$\tilde{\Delta}_i = \Delta_i / \max\left(1, \frac{\|\Delta_i\|_2}{C}\right) + \mathcal{N}(0, \sigma^2 C^2 I) \quad (10)$$

where C is the clipping threshold and σ controls noise magnitude.

Privacy-utility tradeoff analysis (Figure 4): With $\epsilon = 2.1$ – DP (considered strong privacy), TrustLearn maintains 89.7% of non-private accuracy. At $\epsilon = 5.0$, accuracy degradation is only 3.2%.

VI.6 SCALABILITY ANALYSIS

We evaluate scalability across agent populations from 100 to 10,000.

Table 5: Scalability Metrics.

Agents	Round Time (s)	Communication (MB)	Energy (J)	Accuracy
100	12.3 $\pm$ 0.5	8.7 $\pm$ 0.3	19.7 $\pm$ 0.8	0.904 $\pm$ 0.003
500	28.6 $\pm$ 1.2	42.1 $\pm$ 1.5	23.4 $\pm$ 1.0	0.911 $\pm$ 0.003
1,000	51.4 $\pm$ 2.1	83.5 $\pm$ 2.8	28.1 $\pm$ 1.3	0.916 $\pm$ 0.002
5,000	187.2 $\pm$ 7.3	401.2 $\pm$ 12.5	47.8 $\pm$ 2.1	0.921 $\pm$ 0.002
10,000	342.8 $\pm$ 13.6	784.3 $\pm$ 25.4	71.2 $\pm$ 3.4	0.923 $\pm$ 0.002

Source: Authors, (2026).




Figure 4: Privacy-utility tradeoff: Model accuracy vs. epsilon (ε) for different noise levels.

Source: Authors, (2026).

Observations:

- Training time scales sub-linearly due to parallel local training (342.8s for 10K agents vs. 27.8x linear scaling).
- Communication overhead grows linearly but remains manageable (784 MB for 10 K agents per round).
- Energy consumption increases moderately (3.6x for 100x agents), demonstrating efficiency of distributed architecture.
- Model accuracy improves slightly with more agents (+2.1% from 100 to 10 K), indicating effective knowledge aggregation.

VI.7 STATISTICAL VALIDATION

We perform Friedman test [29] to assess statistical significance across methods and datasets, followed by post-hoc Nemenyi test [30] for pairwise comparisons.

Friedman Test Results:

- AUFC metric: $\chi^2(6) = 42.73, p < 0.001$
- Engagement metric: $\chi^2(6) = 39.81, p < 0.001$

Nemenyi Post-hoc: TrustLearn significantly outperforms all baselines at $p < 0.05$ level. Critical difference (CD) at 95% confidence: 1.82.

Effect sizes (Cohen's d) comparing TrustLearn vs. baselines:

- vs. FD: $d = 2.87$ (very large)
- vs. RBA: $d = 2.13$ (very large)
- vs. CDRL: $d = 1.56$ (large)
- vs. FL-noBC: $d = 0.72$ (medium).

VI.8 ENERGY EFFICIENCY IN NZEC CONTEXT

We simulate a NZEC with 1,000 households, each with 1-2 learners and solar energy generation. Figure 5 shows the daily energy consumption profile.




Figure 5: Daily energy consumption profile in simulated NZEC showing adaptive scheduling aligned with solar generation peaks.

Source: Authors, (2026).

TrustLearn's energy-aware scheduling aligns FL rounds with peak solar generation (11 AM - 3 PM), achieving:

- 67% renewable energy utilization for FL operations
- 34% reduction in grid energy consumption vs. non-adaptive scheduling
- Zero-carbon learning operations during 6 hours of daily peak generation.

VI.9 QUALITATIVE ANALYSIS: LEARNED POLICIES

We analyze learned agent policies using saliency mapping to identify critical state features. Top-3 influential features:

1. Current knowledge state (42% importance): Agents heavily weight mastery level when selecting difficulty.
2. Time since last session (28% importance): Spacing effect modeling drives frequency decisions.
3. Recent motivation trend (18% importance): Declining engagement triggers adjustments to re-engage learners.

Policy interpretability is enhanced through decision tree approximation (Fidelity: 87.3%), enabling educators to understand and validate agent recommendations.

VII. DISCUSSION

VII.1 PRINCIPAL FINDINGS

This research establishes TrustLearn as a viable framework for trustworthy adaptive microlearning in resource-constrained environments. Three principal contributions emerge:

1. Empirically-Validated Dosage Optimization: We provide the first large-scale evidence that personalized microlearning dosage significantly impacts retention and motivation. The discovered patterns—shorter, frequent, difficult sessions for advanced learners; longer, spaced, easier sessions for struggling learners—offer actionable insights for educational technology design.
2. End-to-End Trustworthiness: The integration of blockchain, MARL, and FL creates a system where privacy, security, auditability, and performance coexist. Byzantine-resilient aggregation and blockchain-based reputation systems successfully defend against coordination attacks, addressing a critical gap in distributed educational AI.
3. Sustainability Feasibility: The modest 4.3% energy overhead demonstrates that trustworthy, decentralized learning need not compromise sustainability goals. Energy-aware scheduling leverages renewable generation patterns, enabling zero-carbon educational AI operations during peak solar hours.

VII.2 THEORETICAL IMPLICATIONS

VII.2.1 Rethinking Educational Ai Architecture

TrustLearn challenges the centralized paradigm dominating educational AI. Our results suggest that:

- Decentralized learning can outperform centralized approaches (+19.6% vs. CDRL) when properly coordinated via MARL.
- Privacy preservation via FL need not sacrifice personalization; indeed, federated knowledge aggregation enhances individual agent performance.
- Blockchain's role extends beyond credential management to enabling trustworthy AI coordination.

VII.2.2 Cognitive Science Insights

The learned dosage policies align with established cognitive science principles:

- Spacing Effect [31]: High knowledge retention correlates with optimal spacing intervals learned by agents.

- Desirable Difficulties [32]: Advanced learners benefit from challenging content, while struggling learners require scaffolding. Cognitive Load Theory [33]: Session duration adapts to learner capacity, preventing overload. These alignments validate MARL as a viable approach for discovering pedagogically sound strategies.

VII.3 PRACTICAL IMPLICATIONS

VII.3.1 For Educational Technology Developers

1. Adopt Federated Architecture: Privacy-preserving learning is achievable without centralized data collection. This addresses growing regulatory pressures (GDPR, COPPA) and user concerns.
2. Implement Adaptive Dosage: Fixed session parameters underperform personalized approaches by 23 – 31%. Investment in adaptive systems yields measurable engagement improvements.
3. Design for Sustainability: Energy-aware scheduling reduces environmental impact while maintaining performance. This is particularly relevant as AI energy consumption grows.

VII.3.2 For Nzecs And Policymakers

1. Educational Infrastructure Planning: NZECs should provision for distributed learning infrastructure, recognizing its compatibility with renewable energy patterns.
2. Data Sovereignty: TrustLearn enables communities to retain data ownership while benefiting from collaborative learning, supporting digital autonomy.
3. Standards Development: Blockchain-based learning records facilitate credential portability and lifelong learning tracking across platforms.

VII.4 LIMITATIONS AND FUTURE DIRECTIONS

VII.4.1 Current Limitations

1. Simulation vs. Real Deployment: Our NZEC evaluation relies on simulated energy profiles. Real-world deployment would face additional challenges (network unreliability, hardware heterogeneity).
2. Participant Diversity: Datasets primarily cover language learning and standardized assessments. Generalization to other domains (vocational training, professional development) requires validation.
3. Blockchain Scalability: Ethereum's throughput limitations (15 – 30 TPS) constrain system scale. Layer-2 solutions or alternative blockchains may be necessary for massive deployments.
4. Long-Term Retention: Our evaluation window (30 days) captures short-term retention. Multi-year longitudinal studies would assess long-term learning outcomes.
5. Cold Start Problem: New learners lack historical data for personalization. Transfer learning from similar users partially addresses this but requires further research.

VII.4.2 Future Research Directions

1. Multi-Modal Learning: Extend TrustLearn to incorporate visual, auditory, and kinesthetic learning modalities, optimizing modality selection alongside dosage parameters.
2. Social Learning Integration: Incorporate peer collaboration dynamics into MARL formulation, enabling agents to leverage social learning benefits.
3. Explainable AI for Educators: Develop interpretability tools enabling teachers to understand, validate, and override agent recommendations when pedagogically appropriate.
4. Cross-Domain Transfer: Investigate policy transfer across learning domains (e.g., language learning to math) to accelerate personalization for new subjects.
5. Neuroscience-Informed Rewards: Integrate physiological signals (EEG, eye-tracking) to refine reward functions beyond behavioral proxies.
6. Real-World NZEC Pilot: Deploy TrustLearn in an actual Net-Zero Energy Community to validate energy models and assess user acceptance.

VIII. CONCLUSION

This paper introduced TrustLearn, a novel framework integrating blockchain, multi-agent reinforcement learning, and federated learning to address the critical question: What microlearning dosage maximizes knowledge retention and motivation in resource-constrained, privacy-sensitive environments?

Through comprehensive experimentation on three large-scale educational datasets, we demonstrated that:

1. Personalized dosage optimization significantly improves learning outcomes, achieving 23.7% retention gains and 31.2% engagement improvements over fixed approaches.

2. End-to-end trustworthiness is achievable via blockchain-enhanced federated MARL, providing privacy guarantees ($\epsilon = 2.1 - DP$), attack resistance (90% accuracy maintained under 30% Byzantine agents), and auditability.
3. Sustainability and scalability are compatible with trustworthy AI, requiring only 4.3% energy overhead while scaling to 10,000 distributed learners.

TrustLearn represents a paradigm shift toward decentralized, privacy-preserving, and energyaware educational AI. As Net-Zero Energy Communities proliferate and privacy regulations tighten, such frameworks become not merely desirable but essential. Our work provides both theoretical foundations and practical blueprints for building the next generation of adaptive learning systems—systems that respect learner privacy, operate sustainably, and demonstrably enhance educational outcomes. The convergence of educational technology, distributed AI, and sustainable computing opens exciting avenues for future research. We envision a future where learners worldwide benefit from personalized education while retaining data ownership, where AI systems operate transparently and accountably, and where technological progress aligns with environmental sustainability. TrustLearn is a step toward that future.

IX. DATA AVAILABILITY STATEMENT

The datasets analyzed in this study are publicly available: Duolingo SLAM (<https://sharedtask.duolingo.com/>), EdNet (<https://github.com/riid/ednet>), ASSISTments 2009-2010 (<https://sites.google.com/site/assistmentsdata/>). Code and experimental configurations will be released upon publication at <https://github.com/anonymous/trustlearn>.

X. AUTHOR'S CONTRIBUTION

Conceptualization: Mohamed Goismi, Djamel Seghier, Moustafa Maaskri and Mohamed Debbab.

Methodology: Moustafa Maaskri and Mohamed Debbab.

Investigation: Mohamed Goismi.

Discussion of results: Mohamed Goismi, Djamel Seghier, Moustafa Maaskri and Mohamed Debbab.

Writing – Original Draft: Moustafa Maaskri and Mohamed Debbab.

Writing – Review and Editing: Mohamed Debbab.

Resources: Mohamed Goismi, Djamel Seghier, Moustafa Maaskri and Mohamed Debbab.

Supervision: Mohamed Goismi.

Approval of the final text: Mohamed Goismi, Djamel Seghier, Moustafa Maaskri and Mohamed Debbab.

XI. REFERENCES

- [1] G. Barone, A. Buonomano, C. Forzano, G. F. Giuzio, and A. Palombo, "Building to building energy exchange: A review of enabling technologies and applications," *Sustainable Cities and Society*, vol. 62, p. 102454, 2020.
- [2] T. Hug, "Micro learning and narration: exploring possibilities of utilization of narrations and storytelling for the designing of "micro units" and didactical micro-learning arrangements," *Proceedings of the Fourth Media in Transition Conference*, 2005.
- [3] L. Giurgiu and M. Barla, "Microlearning: A new concept in e-learning," in *International Conference on Advanced Learning Technologies*. IEEE, 2017, pp. 320-324.
- [4] K. Zhang, Z. Yang, and T. Başar, "Multi-agent reinforcement learning: A selective overview of theories and algorithms," *Handbook of Reinforcement Learning and Control*, pp. 321-384, 2021.
- [5] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communicationefficient learning of deep networks from decentralized data," in *Artificial Intelligence and Statistics*. PMLR, 2017, pp. 1273-1282.
- [6] Z. Zheng, S. Xie, H. Dai, X. Chen, and H. Wang, "Blockchain challenges and opportunities: A survey," *International Journal of Web and Grid Services*, vol. 14, no. 4, pp. 352-375, 2018.
- [7] G. S. Mohammed, K. Wakil, and S. S. Nawroly, "Microlearning: A modernized education system," *BRAIN. Broad Research in Artificial Intelligence and Neuroscience*, vol. 9, no. 2, pp. 23-35, 2018.
- [8] M. Reddy, K. Ravi, and S. Chandra, "Deep learning based recommendation system for personalized education," in *International Conference on Computing and Network Communications*. IEEE, 2020, pp. 1-6.
- [9] R. Lowe, Y. I. Wu, A. Tamar, J. Harb, O. Pieter Abbeel, and I. Mordatch, "Multi-agent actor-critic for mixed cooperative-competitive environments," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [10] T. Rashid, M. Samvelyan, C. Schroeder, G. Farquhar, J. Foerster, and S. Whiteson, "Qmix: Monotonic value function factorisation for decentralised multi-agent reinforcement learning," in *International Conference on Machine Learning*. PMLR, 2018, pp. 4295-4304.
- [11] D. Hooshyar, M. Pedaste, and Y. Yang, "A multiagent reinforcement learning-based system for game difficulty adaptation," *IEEE Transactions on Learning Technologies*, vol. 13, no. 4, pp. 758-771, 2020.
- [12] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Transactions on Intelligent Systems and Technology*, vol. 10, no. 2, pp. 1-19, 2019.

- [13] Y. Liu, T. Chen, and Q. Yang, "Privacy-preserving federated learning for educational data mining," in International Conference on Educational Data Mining, 2020, pp. 1-10.
- [14] Y. Chen, J. Wang, C. Yu, W. Gao, and X. Qin, "Asynchronous federated learning for geospatial applications," in Joint European Conference on Machine Learning and Knowledge Discovery in Databases. Springer, 2020, pp. 21-37.
- [15] A. Grech and A. F. Camilleri, "Blockchain in education," in JRC Science for Policy Report, European Commission, 2017.
- [16] M. Sharples and J. Domingue, "The blockchain and kudos: A distributed system for educational record, reputation and reward," Adaptive and Adaptable Learning, pp. 490-496, 2016.
- [17] P. Ocheja, B. Flanagan, and H. Ogata, "Blockchain-enabled student profile management and learning analytics," in International Conference on Advanced Learning Technologies. IEEE, 2019, pp. 50-52.
- [18] R. Schwartz, J. Dodge, N. A. Smith, and O. Etzioni, "Green ai," Communications of the *ACM*, vol. 63, no. 12, pp. 54-63, 2020.
- [19] H. Ebbinghaus, "Memory: A contribution to experimental psychology," Annals of Neurosciences, vol. 20, no. 4, p. 155, 1913.
- [20] C. Yu, A. Velu, E. Vinitzky, Y. Wang, A. Bayen, and Y. Wu, "The surprising effectiveness of ppo in cooperative multi-agent games," in Advances in Neural Information Processing Systems, vol. 34, 2021, pp. 24611-24624.
- [21] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, "Practical secure aggregation for privacy-preserving machine learning," in Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, 2017, pp. 1175-1191.
- [22] P. Blanchard, E. M. El Mhamdi, R. Guerraoui, and J. Stainer, "Machine learning with adversaries: Byzantine tolerant gradient descent," Advances in Neural Information Processing Systems, vol. 30, 2017.
- [23] F. Bravo-Marquez, S. Reeves, and M. Ugarte, "Proof of learning: A blockchain consensus mechanism based on machine learning competitions," in International Conference on Distributed Computing Systems Workshops. IEEE, 2021, pp. 119-124.
- [24] B. Settles, C. Brust, E. Gustafson, M. Hagiwara, and N. Madnani, "Second language acquisition modeling," Proceedings of the Thirteenth Workshop on Innovative Use of NLP for Building Educational Applications, pp. 56-65, 2018.
- [25] Y. Choi, Y. Lee, J. Cho, J. Baek, B. Kim, Y. Cha, D. Shin, C. Bae, and J. Heo, "Ednet: A large-scale hierarchical dataset in education," Artificial Intelligence in Education, pp. 69-73, 2020.
- [26] M. Feng, N. Heffernan, and K. Koedinger, "Addressing the assessment challenge with an online system that tutors as it assesses," User Modeling and User-Adapted Interaction, vol. 19, no. 3, pp. 243-266, 2009.
- [27] A. T. Corbett and J. R. Anderson, "Knowledge tracing: Modeling the acquisition of procedural knowledge," User Modeling and User-Adapted Interaction, vol. 4, no. 4, pp. 253-278, 1994.
- [28] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," pp. 308-318, 2016.
- [29] M. Friedman, "The use of ranks to avoid the assumption of normality implicit in the analysis of variance," Journal of the American Statistical Association, vol. 32, no. 200, pp. 675-701, 1937.
- [30] P. Nemenyi, "Distribution-free multiple comparisons," Ph.D. Thesis, Princeton University, 1963.
- [31] N. J. Cepeda, H. Pashler, E. Vul, J. T. Wixted, and D. Rohrer, "Distributed practice in verbal recall tasks: A review and quantitative synthesis," Psychological Bulletin, vol. 132, no. 3, p. 354, 2006.
- [32] R. A. Bjork, "Introducing desirable difficulties for enhancing learning and instruction," Psychology and the Real World: Essays Illustrating Fundamental Contributions to Society, pp. 56-64, 2011.
- [33] J. Sweller, "Cognitive load during problem solving: Effects on learning," Cognitive Science, vol. 12, no. 2, pp. 257-285, 1988.