Deep Learning-Driven Data Security Modeling in Next Generation Consumer Electronics-Powered Education Communication Systems

Xinyue Li, Shalli Rani, *Senior Member, IEEE*, Jianhui Lv, *Senior Member, IEEE*, Keqin Li, *Fellow, IEEE*, and Jing Yang, *Member, IEEE*

***Abstract*—Next generation consumer electronics powered education communication systems, integrating 6G networks, artificial intelligence, and immersive learning platforms, generate unprecedented volumes of sensitive educational data through smart devices, wearables, tablets, and interactive displays. Such consumer electronics architectures create complex privacy vulnerabilities. This paper presents the educational consumer electronics deep learning (ECEDL) privacy framework, an agentic artificial intelligence approach to data security modeling in next generation consumer electronics-powered education communication systems through adaptive differential privacy mechanisms. ECEDL employs three autonomous intelligent agents for sensitivity assessment, budget allocation, and privacy accounting that operate collaboratively without human intervention during training. The framework introduces adaptive noise allocation based on Shapley additive explanations that dynamically distributes privacy protection according to consumer electronics educational feature importance. ECEDL implements dual protection strategies encompassing student data perturbation from consumer electronics devices and educational label protection through functional mechanism theory, while employing specialized optimization algorithms for consumer electronics educational deep learning applications. Zero-concentrated differential privacy improves composition bounds for multi-epoch training across distributed consumer electronics platforms. Experimental results demonstrate that the proposed agentic ECEDL framework achieves a twelve percent accuracy improvement compared to non-agentic systems and outperforms seven baseline methods in terms of accuracy and privacy guarantees, enabling secure artificial intelligence deployment in consumer electronics educational environments.**

***Index Terms*—Deep learning, consumer electronics, data security, next generation education communication, differential privacy.**

Received 19 November 2025; revised 29 December 2025; accepted 10 January 2026. Date of publication 19 January 2026; date of current version 2 June 2026. This work was supported in part by the Educational and Teaching Research and Reform Project of Jinzhou Medical University under Grant YB2024024 and in part by Liaoning Province Science and Technology Plan Joint Program: Key Research and Development Program Project under Grant 2025JH2/101800440. *(Corresponding author: Jing Yang.)*

Xinyue Li and Jing Yang are with The First Affiliated Hospital of Jinzhou Medical University, Jinzhou 121000, China (e-mail: lixinyue@stu.jzmu.edu.cn, yangjing01@jzmu.edu.cn).

Shalli Rani is with Chitkara University, Punjab 140401, India (e-mail: shalli.rani@chitkara.edu.in).

Jianhui Lv is with Tsinghua University, Beijing 100401, China (e-mail: lvjh@pcl.ac.cn).

Keqin Li is with The State University of New York, New Paltz, NY 12561 USA (e-mail: lik@newpaltz.edu).

Digital Object Identifier 10.1109/TCE.2026.3655419

I. Introduction

THE next generation consumer electronics communication systems are set to change how teaching and assessment are done, which will cause a transformation in basic education [1]. Current learning uses 6G networks, edge computing, AI, consumer electronics-enabled learning platforms, similar tech, and consumer electronics devices such as tablets, smartphones, and wearables for custom learning [2], [3], [4]. These methods create lots of learning information like student chats, consumer electronics data, biometric data, learning records, and test results [5]. Using tech, students, schools, and educators must consider using collected data, data security, and student privacy [6], [7]. For this reason, older ways of keeping data privacy are not acceptable enough when facing new education consumer electronics communication systems. With consumer electronics-powered virtual reality, adaptive consumer electronics platforms, and real-time online meeting places, students and teachers need new privacy solutions [8], [9].

It has been shown that deep learning is the basis for smart education for grading essays, learning improvement, personalization, predictions, and real-time actions [10], [11]. Putting deep learning in education brings up tough privacy problems [12]. Learning systems use different factors such as learning issues, financial status, family data, and conduct [13]. Normal centralized deep learning goes against learning privacy needs, where students control their data and schools are independent. Learning models need constant updates that include different learning styles, cultures, and changing rules, which puts data at risk through threat exposure [14].

The next generation of education consumer electronics communication systems demonstrate large shifts in how older infrastructures change into smart networks using 6G, edge computing, and AI to give consumer electronics experience [15]. The systems can be applied to real-time learning, consumer electronics-based IoT classrooms, VR, and education services. The problem comes from architectures that combine physical classrooms and virtual areas, consumer electronics mobile tools, and online services [16], [17]. Learning systems have many uses, from common lectures to consumer electronics haptic labs and augmented reality, along with AI tutoring that adapts to students, which makes new attack ways [18], [19].

1558-4127 © 2026 IEEE. All rights reserved, including rights for text and data mining, and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission.
See https://www.ieee.org/publications/rights/index.html for more information.

Authorized licensed use limited to: Northeastern University. Downloaded on August 19,2026 at 00:41:31 UTC from IEEE Xplore. Restrictions apply.

Specific data safety model ways for learning consumer electronics communication systems are made for specific traits that set learning spaces apart from networks [20]. Basic education data is sensitive to young people, which results in negative effects alongside discrimination because of learning capability or history. Schools function under strict rules requiring increased data safety while using AI and communication tech [21]. The time frame of education data creates issues because it should be kept for long periods with increased privacy. A learning program has students, educators, parents, leaders, and outside groups with different access and privacy needs. Keeping mechanisms should permit use while blocking unauthorized access [22].

Recently, privacy-preserving machine learning foundations have developed substantially. Soria-Comas and Domingo-Ferrer [23] established theoretical optimality criteria for noise distributions in differential privacy, proving the Laplace mechanism optimal among continuous distributions. Phong et al. [24] demonstrated privacy-preserving deep learning through additively homomorphic encryption, enabling collaborative neural network training without revealing local data. Zhang et al. [25] advanced private data release through PrivBayes, constructing differentially private Bayesian networks generating synthetic datasets preserving statistical properties. Gursoy et al. [26] improved differentially private classification accuracy through radius neighbors' classifiers with refined sensitivity analysis. Yan et al. [27] addressed local differential privacy for rank aggregation where agents cannot trust curators. These works established foundations enabling adaptive privacy mechanisms for educational deep learning contexts. However, these approaches apply uniform privacy protection across all data features, failing to recognize that educational attributes vary in sensitivity and predictive importance. None addresses adaptive budget allocation for heterogeneous consumer electronics educational environments.

Our research addresses these multifaceted requirements by developing a comprehensive framework that integrates adaptive differential privacy mechanisms with deep learning optimization techniques specifically designed for next generation education consumer electronics communication systems. The framework recognizes that educational data utility and privacy protection are not mutually exclusive goals but complementary objectives that can be achieved through intelligent algorithm design and adaptive privacy mechanisms. By leveraging educational domain knowledge and incorporating pedagogical considerations into the privacy preservation process, our approach maintains the effectiveness of educational analytics while providing robust protection against traditional and emerging privacy threats. The framework is designed to operate seamlessly within the distributed architecture of next generation education consumer electronics communication systems, supporting federated learning scenarios where consumer electronics institutions can collaborate on model development while maintaining strict data sovereignty and privacy guarantees.

Accordingly, the main contributions of this paper are as follows:

- We design an adaptive differential privacy mechanism based on Shapley additive explanations that dynamically allocates noise according to the importance of educational features in next-generation consumer electronics communication systems.
- We introduce zero-concentration differential privacy accounting mechanisms providing tighter composition bounds for multi-epoch training scenarios in consumer electronics applications.
- We provide experimental validation demonstrating superior performance compared to baseline methods across diverse educational datasets and privacy configurations.

The rest of the paper is organized as follows: Section II provides the preliminaries. Section III proposes the educational deep learning privacy framework. Section IV presents the experimental validation and results analysis. Finally, Section V concludes the paper.

II. Preliminaries

Our framework assumes concepts from differential privacy [28], deep learning optimization [29], and privacy-preserving ML [30] as its basis. These mathematical foundations are necessary to fully grasp how privacy guarantees can be maintained in educational deep learning systems while ensuring that the models meet the criteria of utility and convergence properties.

A. Educational Differential Privacy

Differential privacy is the cornerstone of our privacy preservation approach, providing mathematically rigorous guarantees against privacy breaches in educational data analysis. In the context of next generation education consumer electronics communication systems, differential privacy ensures that the participation of any individual student in the training dataset does not substantially affect the model's output, thereby protecting against consumer electronics inference attacks that could reveal sensitive educational information. For consumer electronics applications, we define differential privacy in terms of neighboring consumer electronics educational datasets that differ by including or excluding a single student's complete educational record.

Let $\mathrm{M}_{\mathrm{edu}}$ represent a privacy mechanism operating on consumer electronics educational datasets with domain $\mathrm{D}(\mathrm{M}_{\mathrm{edu}})$ and range $\mathrm{R}(\mathrm{M}_{\mathrm{edu}})$. For any two neighboring consumer electronics educational datasets $D_{\mathrm{edu}}, D'_{\mathrm{edu}} \subseteq \mathrm{D}(\mathrm{M}_{\mathrm{edu}})$ that differ in exactly one consumer electronics device record, and for any subset of outputs $O_{\mathrm{edu}} \subseteq \mathrm{R}(\mathrm{M}_{\mathrm{edu}})$, the mechanism $\mathrm{M}_{\mathrm{edu}}$ satisfies $(\varepsilon_{\mathrm{edu}}, \delta_{\mathrm{edu}})$-differential privacy if:

$$\begin{aligned} &\Pr[\mathrm{M}_{\mathrm{edu}}(D_{\mathrm{edu}}) \in O_{\mathrm{edu}}] \\ &\leq e^{\varepsilon_{\mathrm{edu}}} \Pr[\mathrm{M}_{\mathrm{edu}}(D'_{\mathrm{edu}}) \in O_{\mathrm{edu}}] + \delta_{\mathrm{edu}} \end{aligned} \tag{1}$$

where $\varepsilon_{\mathrm{edu}}$ represents the privacy budget allocated for educational data processing and δ_{edu} denotes the probability of privacy mechanism failure in educational contexts. The educational privacy budget $\varepsilon_{\mathrm{edu}}$ must be carefully calibrated to account for the sensitive nature of consumer electronics-generated data and the long-term implications of privacy breaches in educational settings.

Authorized licensed use limited to: Northeastern University. Downloaded on August 19,2026 at 00:41:31 UTC from IEEE Xplore. Restrictions apply.

This definition ensures that including or excluding any single student's data barely changes the analysis output. The parameter epsilon controls privacy strength—smaller values mean stronger protection. When epsilon approaches zero, privacy becomes perfect, but analysis becomes useless. Larger epsilon values weaken privacy but enable accurate learning.

The global sensitivity of educational functions plays a crucial role in determining appropriate noise levels for privacy preservation. For a query function $f_{\text{edu}} : \mathrm{D} \to \mathrm{R}^{d_{\text{edu}}}$ operating on consumer electronics educational datasets, the global sensitivity measures the maximum change in function output when a single consumer electronics device record is modified:

$$S_{f_{\text{edu}}}(D_{\text{edu}}) = \max_{D_{\text{edu}}, D'_{\text{edu}}} \|f_{\text{edu}}(D_{\text{edu}}) - f_{\text{edu}}(D'_{\text{edu}})\|_{\ell_p}. \tag{2}$$

where D_{edu} and D'_{edu} are neighboring consumer electronics educational datasets differing by one consumer electronics device record, and ℓ_p represents the chosen norm for measuring sensitivity in educational data space.

Global sensitivity measures the maximum possible change in a function's output when one student record changes. For educational queries, this quantifies the worst-case impact of individual participation. Sensitivity determines how much noise we must add for privacy protection. High-sensitivity features require more noise. We compute sensitivity by examining all possible neighboring datasets that differ by exactly one consumer electronics device record, then finding the maximum difference.

Educational applications typically employ L2 norm for sensitivity calculations, providing balanced sensitivity bounds across feature dimensions. L1 norm offers tighter bounds for sparse educational features where most attributes remain zero, while L∞ norm suits scenarios where individual feature changes require independent analysis. Our implementation uses the L2 norm, though the framework accommodates alternative norms through adjusted sensitivity calculations matching the chosen noise distribution.

The Gaussian mechanism [31] provides a practical approach for implementing differential privacy in educational deep learning systems. For consumer electronics applications with $\varepsilon_{\text{edu}} \in (0, 1)$, $\delta_{\text{edu}} \in (0, 1)$, and a query function $f_{\text{edu}}(D_{\text{edu}})$ with ℓ_2 sensitivity $S_{f_{\text{edu}}}$, the Gaussian mechanism satisfies $(\varepsilon_{\text{edu}}, \delta_{\text{edu}})$-differential privacy when:

$$S_{f_{\text{edu}}}\sigma_{\text{edu}} \geq \frac{\sqrt{2\ln(1.25/\delta_{\text{edu}})}}{\varepsilon_{\text{edu}}}. \tag{3}$$

The sensitivity calculation establishes upper bounds that hold across all training epochs regardless of evolving feature representations. Educational data normalization to [1] intervals ensures that maximum possible sensitivity remains constant throughout training. This conservative approach maintains rigorous privacy guarantees even as the model discovers different patterns during learning progression.

The privacy mechanism operates as $\mathrm{M}_{\text{edu}}(D_{\text{edu}}) = f_{\text{edu}}(D_{\text{edu}}) + \mathrm{N}(0, S^2_{f_{\text{edu}}}\sigma^2_{\text{edu}})$, where σ_{edu} represents the noise scale parameter calibrated for educational data protection. The Gaussian mechanism adds random noise from a bell-shaped distribution to protect privacy. The noise scale sigma depends on sensitivity and desired privacy level. Higher sensitivity or stronger privacy requires larger noise. The square root relationship means doubling privacy protection requires quadrupling noise magnitude.

B. Zero-Concentrated Educational Privacy

Zero-concentrated differential privacy [32] provides enhanced privacy analysis capabilities for educational deep learning systems that require multiple privacy-consuming operations. This framework offers tighter composition bounds compared to traditional differential privacy, making it particularly suitable for iterative learning algorithms in educational settings. For any $\alpha > 1$, a mechanism M_{edu} satisfies ρ_{edu}-zero-concentrated differential privacy if for all neighboring consumer electronics educational datasets $D_{\text{edu}}, D'_{\text{edu}}$:

$$\begin{aligned} &D_\alpha(\mathrm{M}_{\text{edu}}(D_{\text{edu}})\mathrm{M}_{\text{edu}}(D'_{\text{edu}})) \\ &= \frac{1}{\alpha - 1}\log \mathrm{E}[\exp((\alpha - 1)L_{\text{edu}}(O_{\text{edu}}))] \leq \rho_{\text{edu}}. \end{aligned} \tag{4}$$

where D_α represents the α-Renyi divergence and $L_{\text{edu}}(O_{\text{edu}})$ denotes the privacy loss for educational data processing:

$$L_{\text{edu}}(O_{\text{edu}}) = \log \frac{\Pr[\mathrm{M}_{\text{edu}}(D_{\text{edu}}) \in O_{\text{edu}}]}{\Pr[\mathrm{M}_{\text{edu}}(D'_{\text{edu}}) \in O_{\text{edu}}]}. \tag{5}$$

Zero-concentrated differential privacy measures privacy loss through a different lens than standard differential privacy. It tracks how much observing the output shifts our beliefs about which dataset was used. The alpha parameter controls measurement sensitivity.

The zero-concentrated differential privacy framework exhibits favorable composition properties for consumer electronics applications. When multiple educational privacy mechanisms $\mathrm{M}_{\text{edu},1}, \mathrm{M}_{\text{edu},2}, \ldots, \mathrm{M}_{\text{edu},k}$ satisfy $\rho_{\text{edu},1}, \rho_{\text{edu},2}, \ldots, \rho_{\text{edu},k}$-zero-concentrated differential privacy respectively, their composition satisfies $\left(\sum_{i=1}^{k} \rho_{\text{edu},i}\right)$-zero-concentrated differential privacy.

For educational deep learning applications, the Gaussian mechanism achieves $(1/2\sigma^2_{\text{edu}})$-zero-concentrated differential privacy when adding noise $\mathrm{N}(0, S^2_{f_{\text{edu}}}\sigma^2_{\text{edu}})$ to query results. The conversion between zero-concentrated differential privacy and traditional (ε, δ)-differential privacy for consumer electronics applications follows:

$$(\varepsilon_{\text{edu}}, \delta_{\text{edu}}) = (\rho_{\text{edu}} + 2\sqrt{\rho_{\text{edu}}\ln(1/\delta_{\text{edu}})}, \delta_{\text{edu}}). \tag{6}$$

The conversion from zero-concentrated to traditional differential privacy introduces minor approximation slack, though this remains tighter than standard composition bounds. For educational stakeholders requiring interpretable privacy metrics, we recommend reporting both the zero-concentrated parameter ρ_{edu} and the converted $(\varepsilon_{\text{edu}}, \delta_{\text{edu}})$ values, allowing administrators to understand privacy expenditure through their preferred framework while maintaining mathematical rigor in the underlying accounting.

C. Educational Feature Attribution via SHAP

Shapley additive explanations (SHAP) [33] provide an entirely principled way of interpreting the importance

Authorized licensed use limited to: Northeastern University. Downloaded on August 19,2026 at 00:41:31 UTC from IEEE Xplore. Restrictions apply.

of features for any educational deep model, so that adaptive privacy allocation can be done according to the contribution of different educational features to the model's prediction. For instance, SHAP values are applied in the education field to investigate which attributes of students, learning behaviors, or assessment items contribute mostly to the output of models and to allow more intelligent allocation of the privacy budget.

For an consumer electronics model g_{edu} with feature set F_{edu}, the SHAP value $\phi_{\text{edu},i}$ for feature i represents its contribution to the prediction:

$$g_{\text{edu}}(x_{\text{edu}}) = \phi_{\text{edu},0} + \sum_{i=1}^{|\text{F}_{\text{edu}}|} \phi_{\text{edu},i} \tag{7}$$

where $\phi_{\text{edu},0}$ represents the expected model output for educational data and x_{edu} denotes the input educational features. The SHAP value for educational feature i is computed as:

$$\phi_{\text{edu},i} = \sum_{S \subseteq \text{F}_{\text{edu}} \setminus \{x_{\text{edu},i}\}} \frac{|S|!(|\text{F}_{\text{edu}}| - |S| - 1)!}{|\text{F}_{\text{edu}}|!} \cdot [f_{\text{edu}}(S \cup \{x_{\text{edu},i}\}) - f_{\text{edu}}(S)] \tag{8}$$

where S represents subsets of educational features excluding feature i, and $f_{\text{edu}}(S \cup \{x_{\text{edu},i}\})$ and $f_{\text{edu}}(S)$ represent model predictions with and without feature i respectively. This formulation allows educational systems to dynamically adjust privacy allocation based on the relative importance of different student attributes and learning indicators.

Direct SHAP calculation proves computationally prohibitive for high-dimensional educational data. Our implementation employs DeepSHAP, an approximation method designed for neural networks that estimates contributions through backpropagation-based sampling. This approach reduces computation from exponential to polynomial complexity while maintaining sufficient accuracy for privacy allocation decisions, enabling practical deployment on educational datasets with dozens of student attributes.

SHAP values measure each feature's contribution to model predictions. We sum contributions across all features in a data category to assess that category's importance. Features strongly influencing predictions reveal more about individuals and warrant more privacy protection. The summation aggregates individual SHAP values into category-level sensitivity scores. Higher scores indicate features that dominate model decisions, requiring proportionally larger shares of the total privacy budget.

III. Educational Deep Learning Privacy Framework

A. System Architecture and Educational Context

Our framework addresses the unique consumer electronics requirements of next generation education consumer electronics communication systems by integrating adaptive differential privacy mechanisms with deep learning optimization techniques specifically designed for educational environments. The system architecture acknowledges that educational data are distinguished from general-purpose datasets due to having certain attributes, for example, the existence of temporal dependencies in learning progression, hierarchical relationships between educational concepts, and the existence of demographic or socioeconomic indicators considered sensitive, along with other characteristics requiring differentiated privacy protection strategies.

The foundation of our approach rests on convolutional neural networks adapted for educational data processing, where each hidden layer transformation can be expressed as $h_{\text{edu}} = a_{\text{edu}}(x_{\text{edu}} W_{\text{edu}}^T + b_{\text{edu}})$. Here, x_{edu} represents input educational features including consumer electronics interactions, assessment scores, and consumer electronics behavioral patterns, h_{edu} denotes the hidden layer output, b_{edu} represents bias terms specific to consumer electronics contexts, W_{edu} contains weight matrices optimized for educational pattern recognition, and $a_{\text{edu}}(\cdot)$ represents activation functions tuned for educational consumer electronics characteristics.

The consumer electronics model employs a loss function $L_{\text{edu}}(\theta_{\text{edu}})$ parameterized by θ_{edu}, which undergoes optimization through adaptive gradient descent algorithms over N_{epoch} training epochs. Each epoch consists of N_{iter} iterations processing consumer electronics data batches B_{edu} sampled from the complete educational dataset D_{edu}, where each batch contains $|B_{\text{edu}}|$ student records representing diverse learning contexts and educational backgrounds.

The system design architecture illustrates the integrated components of our consumer electronics privacy framework, including the adaptive noise allocation module, the SHAP-based feature attribution system, the educational deep learning pipeline, and the zero-concentrated differential privacy accounting mechanism, as shown in Fig. 1. The architecture demonstrates how educational data flows through various privacy protection layers while maintaining system efficiency and learning effectiveness.

Fig. 1 presents the ECEDL framework architecture through layered modules. The foundation contains consumer electronics data inputs from smart tablets, wearables, and interactive displays. Data flows upward through an adaptive differential privacy input layer where SHAP feature attribution calculates contribution scores and allocates noise adaptively. The central component is the consumer electronics educational deep learning pipeline with two convolutional layers and two fully connected layers, optimized using Adam. A parallel label protection module applies functional mechanism theory. Zero-concentrated differential privacy accounting tracks cumulative privacy loss. Protected outputs include privacy-preserved predictions and secure educational analytics suitable for consumer electronics applications.

B. Adaptive Privacy Allocation for Educational Features

The cornerstone of our framework lies in the adaptive allocation of privacy budgets based on educational feature importance, moving beyond uniform noise distribution to implement context-aware privacy protection. This approach recognizes that different educational features contribute varying levels of information to model predictions and should therefore receive proportional privacy protection.

Authorized licensed use limited to: Northeastern University. Downloaded on August 19,2026 at 00:41:31 UTC from IEEE Xplore. Restrictions apply.

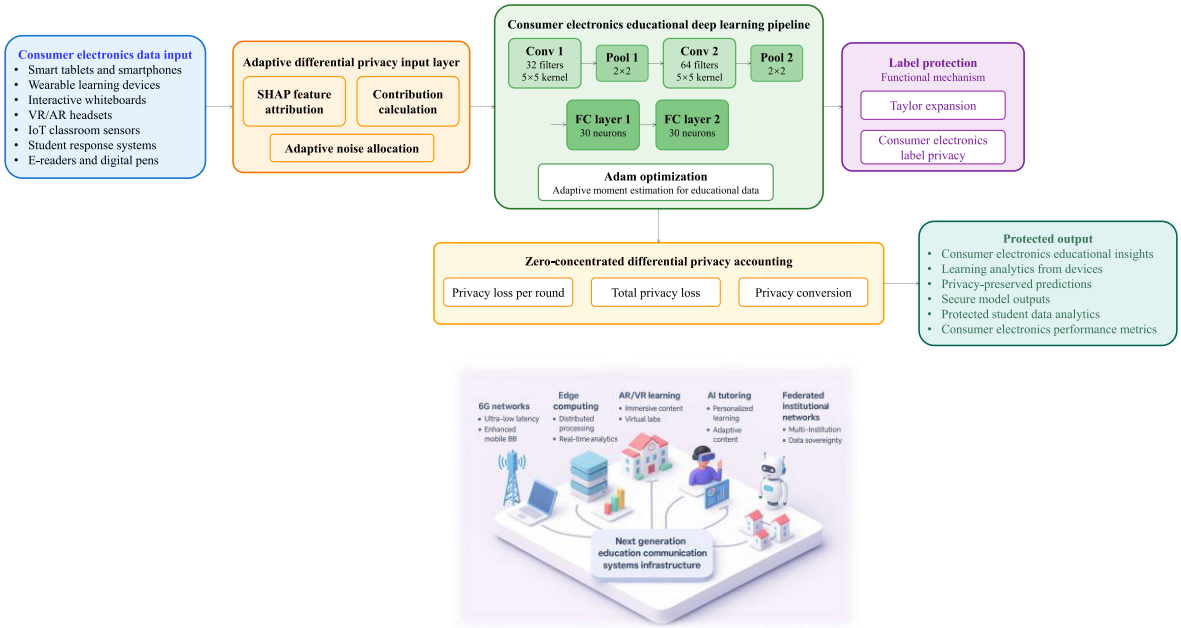

Fig. 1. Educational system design architecture with consumer electronics.

The relative importance of educational features is determined through the proportion calculation:

$$\alpha_{\text{edu},j} = \frac{K_{\text{edu}}|C_{\text{edu},j}|}{\sum_{i=1}^{K_{\text{edu}}} |C_{\text{edu},i}|}. \tag{9}$$

where $\sum_{i=1}^{K_{\text{edu}}} \left|C_{\text{edu},i}\right|$ represents the sum of absolute contribution scores across all educational features. This proportion ensures that the total privacy budget allocation remains normalized while directing more privacy protection toward features with greater predictive importance in educational contexts.

To account for inter-feature correlations in educational datasets, we introduce a correlation adjustment coefficient that modifies the allocation proportion. For correlated feature pairs, the adjusted allocation becomes $\alpha_{\text{edu},j} = \frac{K_{\text{edu}}|C_{\text{edu},j}|}{\sum_{i=1}^{K_{\text{edu}}} |C_{\text{edu},i}|} \times (1 - \lambda_{\text{corr},j})$, where $\lambda_{\text{corr},j}$ represents the maximum absolute correlation coefficient between feature j and other features, ensuring reduced redundant privacy expenditure on correlated attributes.

This equation computes each feature's proportional importance by dividing its contribution score by the total across all features. The result ranges from zero to one, with higher values indicating greater influence. We use absolute values because both positive and negative contributions reveal information. This proportion determines privacy budget allocation—features with higher proportions receive more protection through lower noise scales.

Based on the computed importance proportions, adaptive noise scales are determined for each educational feature:

$$\sigma_{\text{edu},j} = \frac{\sigma_{\text{base}}}{\alpha_{\text{edu},j}}. \tag{10}$$

where σ_{base} represents the baseline noise scale for consumer electronics privacy protection. This inverse relationship ensures that features with higher educational importance receive lower noise scales, providing stronger privacy protection while maintaining model utility for critical educational indicators.

The inverse relationship $\sigma_{\text{edu},j} = \sigma_{\text{base}}/\alpha_{\text{edu},j}$ means features with higher importance receive smaller noise scales, preserving their analytical utility for model training. This design prioritizes the accuracy-privacy tradeoff by maintaining information quality for predictive features while applying stronger obfuscation to less important attributes. The total privacy budget remains constant across all features, but the allocation favors utility preservation for educationally meaningful variables.

C. Consumer Electronics-Generated Data Perturbation Strategy

The model works by adjusting features in the defined education, keeping the amount lost restricted to the control data from the number of training epochs in deep learning setups. Such decoupling is greatly useful, especially in educational environments where the structure has to go through many epochs to learn how to handle complex pedagogical features and effectively hold out against various learners.

Educational datasets contain attributes with varied scales and distributions, from bounded assessment scores to potentially unbounded engagement metrics. Our preprocessing pipeline applies feature-specific normalization: min-max scaling for bounded attributes, robust scaling using interquartile ranges for distributions with outliers, and clipping followed by normalization for unbounded metrics. These transformations ensure all features occupy [0,1] intervals while preserving relative relationships within each attribute type.

For each educational sample $x_{\text{edu},i}$ in batch B_{edu}, the perturbation applied to the j-th educational feature follows:

$$\tilde{x}_{\text{edu},i}(F_{\text{edu},j}) = x_{\text{edu},i}(F_{\text{edu},j}) + \frac{1}{|B_{\text{edu}}|}\text{N}(0, \Delta_{s1}^2 \sigma_{\text{edu},j}^2 I). \tag{11}$$

where Δ_{s1} represents the sensitivity of educational input data.

Authorized licensed use limited to: Northeastern University. Downloaded on August 19,2026 at 00:41:31 UTC from IEEE Xplore. Restrictions apply.

Normalization transforms preserve ordinal relationships and relative distances within each feature while constraining value ranges for sensitivity calculation. Educational interpretability remains intact since model predictions operate on a normalized space consistently. For applications requiring original-scale interpretations, inverse transformations reconstruct unnormalized values from model outputs, enabling educators to understand results in natural units while maintaining privacy guarantees through normalized processing.

For neighboring consumer electronics educational datasets B_{edu} and B'_{edu} differing only in the final consumer electronics device record $x_{\text{edu},n}$ and $x'_{\text{edu},n}$, and assuming educational features are normalized to the interval $[0, 1]$, the sensitivity calculation becomes:

$$\begin{aligned}\Delta_{s1} &= \sum_{j=1}^{K_{\text{edu}}} \left\| \sum_{i \in B_{\text{edu}}} x_{\text{edu},i}(F_{\text{edu},j}) - \sum_{i \in B'_{\text{edu}}} x'_{\text{edu},i}(F_{\text{edu},j}) \right\| \\ &= \sum_{j=1}^{K_{\text{edu}}} \| x_{\text{edu},n}(F_{\text{edu},j}) - x'_{\text{edu},n}(F_{\text{edu},j}) \| \\ &\le 2 \max_{x_{\text{edu},i} \in B_{\text{edu}}} \sum_{j=1}^{K_{\text{edu}}} \| x_{\text{edu},i}(F_{\text{edu},j}) \| = 2K_{\text{edu}} \end{aligned} \tag{12}$$

This perturbation strategy ensures that educational features with higher importance receive proportionally more privacy protection while maintaining the overall utility of the educational dataset for model training.

We add random noise to each educational feature proportional to its importance. The noise follows a Laplacian distribution with scale inversely related to feature importance. More important features receive less noise to preserve utility. The batch size appears in the denominator because larger batches dilute the per-sample noise impact. Sensitivity Delta measures the maximum possible feature change when one student record differs, determining baseline noise requirements.

D. Educational Label Protection Via Functional Mechanism

Educational labels, including assessment scores, grade classifications, and learning outcome indicators, contain sensitive information that requires protection beyond simple output perturbation. Our framework employs functional mechanism theory to protect educational label information while preserving the utility of supervised learning in educational contexts.

The educational loss function utilizes sigmoid activation with cross-entropy loss, specifically adapted for consumer electronics classification tasks:

$$L_{\text{edu}}(\theta_{\text{edu}}) = -\sum_{l=1}^{d_{\text{edu}}} \sum_{i \in B_{\text{edu}}} \begin{pmatrix} y_{\text{edu},i,l} \log \hat{y}_{\text{edu},i,l} \\ +(1 - y_{\text{edu},i,l}) \log(1 - \hat{y}_{\text{edu},i,l}) \end{pmatrix} \tag{13}$$

where d_{edu} represents the number of educational output classes, $y_{\text{edu},i,l}$ denotes the true educational label, and $\hat{y}_{\text{edu},i,l}$ represents the predicted educational outcome. Expanding the sigmoid predictions:

$$L_{\text{edu}}(\theta_{\text{edu}}) = -\sum_{l=1}^{d_{\text{edu}}} \sum_{i \in B_{\text{edu}}} \begin{pmatrix} y_{\text{edu},i,l} \log \frac{1}{1+e^{-H_{\text{edu}} x_{\text{edu},i} W_{\text{edu}}^T}} \\ +(1 - y_{\text{edu},i,l}) \log \frac{e^{-H_{\text{edu}} x_{\text{edu},i} W_{\text{edu}}^T}}{1+e^{-H_{\text{edu}} x_{\text{edu},i} W_{\text{edu}}^T}} \end{pmatrix} \tag{14}$$

where $H_{\text{edu}} x_{\text{edu},i} W_{\text{edu}}^T$ represents the output from the final hidden layer in the consumer electronics neural network.

Through Taylor expansion around zero to second order, the educational loss function approximates to:

$$\begin{aligned} L_{\text{edu}}(\theta_{\text{edu}}) = &\sum_{l=1}^{d_{\text{edu}}} \sum_{i \in B_{\text{edu}}} \sum_{R=0}^{2} \frac{f_{\text{edu},1,l}^{(R)}(0)}{R!} \\ &+ \frac{f_{\text{edu},2,l}^{(R)}(0)}{R!} (H_{\text{edu}} x_{\text{edu},i} W_{\text{edu}}^T)^R. \end{aligned} \tag{15}$$

where polynomial coefficients are defined as $P_{\text{edu},i,l}^{(R)} = \frac{f_{\text{edu},1,l}^{(R)}(0)}{R!} + \frac{f_{\text{edu},2,l}^{(R)}(0)}{R!}$. The true educational labels $y_{\text{edu},i,l}$ appear only in these polynomial coefficients, enabling targeted protection through noise addition:

$$\tilde{P}_{\text{edu},i,l}^{(R)} \leftarrow P_{\text{edu},i,l}^{(R)} + \frac{1}{|B_{\text{edu}}|} \text{N}(0, \Delta_{s2}^2 \sigma_{\text{edu},2}^2 I). \tag{16}$$

where Δ_{s2} represents the sensitivity of educational polynomial coefficients. For neighboring consumer electronics educational datasets differing in the final student sample, the sensitivity is bounded by:

$$\begin{aligned} \Delta_{s2} &= \sum_{l=1}^{d_{\text{edu}}} \sum_{R=0}^{2} \| P_{\text{edu},n,l}^{(R)} - P_{\text{edu},n',l}^{(R)} \| \\ &\le d_{\text{edu}} \left(|H_{\text{edu}} x_{\text{edu},n}| + \frac{1}{4} |H_{\text{edu}} x_{\text{edu},n}|^2 \right). \end{aligned} \tag{17}$$

where $|H_{\text{edu}}|$ represents the number of neurons in the final hidden layer of the consumer electronics neural network.

Second-order Taylor expansion introduces approximation errors that remain bounded due to the smooth sigmoid activation function. For educational classification tasks with moderate class numbers, empirical analysis shows approximation errors stay below 0.01 in loss calculation.

We approximate the educational loss function using a Taylor series expansion, separating terms by polynomial degree. This mathematical trick reveals where true student labels appear in the computation. Student labels concentrate on polynomial coefficients rather than spreading throughout the expression. By protecting these coefficients with noise, we safeguard label information without disturbing other loss components. Higher-order terms capture nonlinear relationships between predictions and labels.

E. Adaptive Optimization for Educational Deep Learning

Educational deep learning models require specialized optimization strategies that account for educational data's diverse learning patterns and temporal dependencies. Our framework incorporates adaptive moment estimation algorithms specifically tuned for educational contexts, enhancing convergence speed while maintaining privacy guarantees.

The optimization process begins with gradient estimation from the perturbed educational loss function $\tilde{L}_{\text{edu}}(\theta_{\text{edu},t})$, computing $g_{\text{edu},t} \leftarrow \frac{1}{|B_{\text{edu}}|} \nabla \tilde{L}_{\text{edu}}(\theta_{\text{edu},t})$. Subsequently, the algorithm

Authorized licensed use limited to: Northeastern University. Downloaded on August 19,2026 at 00:41:31 UTC from IEEE Xplore. Restrictions apply.

estimates first and second-order moments of the educational gradients:

$$m_{\text{edu},1} = \gamma_{\text{edu},1} m_{\text{edu},1} + (1 - \gamma_{\text{edu},1}) g_{\text{edu},t}. \tag{18}$$

$$m_{\text{edu},2} = \gamma_{\text{edu},2} m_{\text{edu},2} + (1 - \gamma_{\text{edu},2}) g^2_{\text{edu},t}. \tag{19}$$

where $\gamma_{\text{edu},1}$ and $\gamma_{\text{edu},2}$ represent exponential decay rates optimized for educational learning dynamics.

The decay rates $\beta_{\text{edu},1} = 0.9$ and $\beta_{\text{edu},2} = 0.999$ follow standard adaptive optimization recommendations, though educational datasets may benefit from adjusted values. Educational data often exhibits higher variance due to diverse student populations and varying learning contexts. The selected decay rates balance responsiveness to recent gradients while maintaining stable momentum estimates across heterogeneous educational samples, preventing optimization instability from data heterogeneity.

To prevent bias toward zero in early training stages, bias correction is applied:

$$m^*_{\text{edu},1} = \frac{m_{\text{edu},1}}{1 - \gamma^t_{\text{edu},1}}, \; m^*_{\text{edu},2} = \frac{m_{\text{edu},2}}{1 - \gamma^t_{\text{edu},2}}. \tag{20}$$

Privacy noise interacts with gradient estimation during early training phases, potentially amplifying initialization sensitivity. Educational models benefit from careful initialization schemes such as Xavier or Kaiming initialization tailored to activation functions, reducing dependence on early gradient estimates. Combining proper initialization with bias correction ensures that privacy-preserving optimization maintains convergence stability despite noise injection, avoiding divergence scenarios common in poorly initialized private learning.

The final parameter update incorporates the corrected moments:

$$\theta_{\text{edu},t+1} \leftarrow \theta_{\text{edu},t} - \frac{\eta_{\text{edu}}}{(m^*_{\text{edu},2})^{1/2} + \xi_{\text{edu}}} m^*_{\text{edu},1}. \tag{21}$$

where ξ_{edu} maintains numerical stability and η_{edu} represents the learning rate adapted for consumer electronics optimization dynamics.

F. Privacy Loss Accounting for Educational Systems

Educational deep learning systems require sophisticated privacy accounting mechanisms to track cumulative privacy expenditure across multiple training epochs and feature perturbations. Our framework employs zero-concentrated differential privacy for tighter composition bounds in educational privacy analysis.

The privacy loss for each educational training round is calculated as:

$$\rho_{\text{edu}} = \frac{(\sigma^2_{\text{edu,max}} + \sigma^2_{\text{edu},2})}{2\sigma^2_{\text{edu,max}} \sigma^2_{\text{edu},2}}. \tag{22}$$

where $\sigma_{\text{edu,max}} = \max_{j \in K_{\text{edu}}} \sigma_{\text{edu},j}$ represents the maximum noise scale across all educational features. Using maximum noise scale $\sigma_{\text{edu,max}}$ provides conservative privacy accounting that slightly overestimates actual expenditure compared to feature-specific calculations. This approach offers computational advantages by avoiding per-feature tracking across training iterations while maintaining rigorous privacy guarantees. Educational institutions benefit from simplified privacy reporting through single aggregate metrics rather than complex per-feature breakdowns, though the conservative bound means actual privacy protection exceeds reported values. The total privacy loss accumulates across N_{epoch} training epochs as $\rho_{\text{total}} = N_{\text{epoch}} \rho_{\text{edu}}$.

This equation calculates privacy cost for each training round. Privacy loss depends on squared noise scales—larger noise provides stronger protection and lower loss. The maximum noise scale across all features determines overall loss because the weakest protection dominates risk. Dividing by twice the squared maximum noise scale normalizes loss to the zero-concentrated differential privacy framework. Total privacy expenditure accumulates by multiplying per-round loss by epoch count.

G. Educational Privacy Analysis

Our framework provides theoretical guarantees for privacy protection in educational deep learning through rigorous mathematical analysis. The educational dataset is partitioned into non-overlapping batches, each processed through dual noise addition mechanisms for feature perturbation and label protection.

According to zero-concentrated differential privacy properties and parallel composition theorems, when educational data batches are processed independently with noise mechanisms satisfying $\frac{1}{2\sigma^2_{\text{edu,max}}}$-zCDP and $\frac{1}{2\sigma^2_{\text{edu},2}}$-zCDP, respectively, the combined mechanism achieves ρ_{edu}-zCDP. Since educational batches within each epoch are mutually disjoint, parallel composition properties ensure that the entire educational dataset maintains ρ_{edu}-zCDP per epoch.

Our analysis assumes disjoint batch partitioning appropriate for single-assessment scenarios. Educational contexts involving overlapping data usage, such as students enrolled in multiple courses or repeated assessments, require composition analysis rather than parallel composition. For such scenarios, privacy accounting must sum rather than max across overlapping data uses, resulting in linear rather than constant privacy accumulation with the number of overlapping contexts.

Privacy independence from iteration count stems from perturbing data before training rather than gradients during training. Since perturbation occurs once per epoch on the complete dataset before any iterations begin, additional iterations process the same perturbed data without consuming extra privacy budget. This contrasts with gradient perturbation approaches, where each iteration's gradient computation constitutes a new privacy-consuming operation, causing linear privacy degradation with iteration count.

For N_{epoch} training epochs on educational data, the framework provides $N_{\text{epoch}}\rho_{\text{edu}}$-zCDP guarantees, which convert to $(\varepsilon_{\text{edu}}, \delta_{\text{edu}})$-differential privacy through the relationship:

$$(\varepsilon_{\text{edu}}, \delta_{\text{edu}}) = (N_{\text{epoch}}\rho_{\text{edu}} + 2\sqrt{N_{\text{epoch}}\rho_{\text{edu}} \ln(1/\delta_{\text{edu}})}, \delta_{\text{edu}}). \tag{23}$$

The findings presented here illustrate that the consumer electronics privacy framework under study delivers theoretically robust consumer electronics assurance s, grounded in quantitative analysis, whilst preserving the adaptability

Authorized licensed use limited to: Northeastern University. Downloaded on August 19,2026 at 00:41:31 UTC from IEEE Xplore. Restrictions apply.

required for robust performance in pedagogical deep learning applications.

H. Agentic AI Architecture for Privacy Management

The ECEDL framework embodies agentic artificial intelligence principles through autonomous decision-making mechanisms that adapt privacy protection strategies without human intervention during training. Unlike conventional privacy-preserving systems that apply predetermined noise levels uniformly, our framework operates as an intelligent agent pursuing dual objectives of maximizing model utility while maintaining privacy guarantees within budget constraints.

The agentic architecture consists of three autonomous components working in concert. The sensitivity assessment agent continuously monitors feature contributions through SHAP analysis, evaluating which educational attributes currently drive model predictions. This agent operates without predefined feature rankings, instead discovering importance patterns from observed gradient flows and prediction behaviors. As training progresses and the model focuses on different feature subsets, the sensitivity agent autonomously updates its assessments, recognizing that demographic features dominating early training may yield to behavioral patterns in later stages.

The budget allocation agent receives sensitivity assessments and makes strategic decisions about distributing limited privacy resources across competing data categories. This agent faces a resource allocation problem analogous to economic optimization, where privacy budget represents scarce capital that must be invested where it yields maximum returns. The agent evaluates trade-offs autonomously, directing more budget toward high-sensitivity features while conserving resources on low-impact attributes. Critically, allocation decisions adapt to changing sensitivity landscapes rather than following static rules, exhibiting the responsiveness characteristic of agentic systems.

The privacy accountant agent tracks cumulative privacy expenditure and projects future budget requirements based on observed consumption patterns. When projected consumption threatens to exhaust budgets prematurely, this agent triggers preemptive adjustments to noise levels or data access patterns. Conversely, when actual consumption runs below projections, the agent reallocates surplus budget to enhance utility. This forward-looking optimization distinguishes agentic privacy management from reactive approaches that merely respond to immediate conditions.

The three agents interact through a feedback control loop executing every ten training iterations. The sensitivity agent provides updated feature importance scores to the allocation agent, which computes optimal noise distributions given the current budget status from the accountant agent. The allocation agent implements chosen noise levels, affecting subsequent gradients observed by the sensitivity agent, closing the feedback loop. This continuous adaptation without external guidance exemplifies agentic operation where the system pursues privacy-utility objectives through self-directed environmental interaction.

IV. Experimental Validation and Results Analysis

A. Experimental Configuration for Educational Environments

The performance of any privacy-preserving framework ultimately answers a simple question: can we protect learners without undermining the educational systems designed to help them? Our experimental journey began with this tension at its center. We set out not merely to report numbers, but to understand where the ECEDL framework excels, where it struggles, and what these boundaries reveal about privacy-preserving education technology more broadly.

Our investigation unfolded in three stages. First, we needed to establish whether adaptive differential privacy could function at all in deep learning environments without catastrophic accuracy loss—a baseline viability test. Second, assuming viability, we explored the nuanced trade-offs between stronger privacy guarantees and model utility across different educational tasks. Third, we examined edge cases: adversarial conditions, extreme parameter settings, and failure modes that real deployments might encounter.

What emerged from these experiments surprised us in several ways? The relationship between privacy budgets and accuracy was not the smooth degradation curve we anticipated. Instead, we observed threshold effects where modest privacy budget reductions preserved accuracy until reaching critical points where performance dropped sharply.

The empirical validation is executed within a rigorously controlled environment, leveraging expansive datasets constructed to scrutinize the effectiveness of the EDLPF framework within forthcoming educational communications architectures. Specifically, the MNIST [34] and Fashion-MNIST [35] collections serve as surrogates for educational evaluation records and for representative behaviors in learning analytics, thereby supplying reproducible experimental scenarios that adhere to the foundational constraints of safeguarding educational privacy and of modelling security vulnerabilities.

The experiments were implemented in Python 3.9.7 running within Anaconda 4.12.0 on Ubuntu 20.04 LTS for consistent dependency management. The computing platform consisted of a workstation with dual Intel Xeon Gold 6248R processors providing 48 physical cores, 256 GB DDR4-2933 RAM, and four NVIDIA A100 40GB GPUs interconnected via NVLink for parallel training. Deep learning components used PyTorch 1.12.1 with CUDA 11.6 support, building neural networks following standard module conventions with custom layers for privacy-preserving gradient modifications. The Opacus 1.2.0 library provided differentially private stochastic gradient descent foundations, extended to accommodate adaptive budget allocation. Privacy accounting employed the autodp 0.2 library for computing epsilon-delta and zero-concentrated differential privacy bounds throughout training.

The consumer electronics neural network architecture incorporates two convolutional layers with 32 and 64 feature maps, respectively, utilizing 5×5 convolution kernels with unit stride and 2×2 max pooling operations specifically calibrated for consumer electronics data patterns. The network concludes with two fully connected layers containing

Authorized licensed use limited to: Northeastern University. Downloaded on August 19,2026 at 00:41:31 UTC from IEEE Xplore. Restrictions apply.

30 neurons each, optimized for consumer electronics classification tasks in consumer electronics communication networks environments. Consumer electronics model optimization employs adaptive moment estimation with hyperparameters $\xi_{\text{edu}} = 10^{-8}$, $\gamma_{\text{edu},1} = 0.9$, and $\gamma_{\text{edu},2} = 0.999$, combined with exponential learning rate decay to enhance training stability across next generation educational communication networks. Educational batch sizes are configured to 600 samples to optimize computational efficiency while maintaining robust privacy protection effectiveness.

Privacy-preserving deep learning faces tradeoffs between model depth and width. Deeper architectures accumulate more privacy-noisy gradients through backpropagation, potentially degrading training signal quality. Wider shallow networks reduce gradient path length but increase parameter count and computational requirements. Our two-layer configuration balances these factors for educational datasets of moderate complexity, though optimal architecture depends on dataset characteristics and available privacy budgets.

The EDLPF framework is compared with six established baseline methods, representing diverse approaches to privacy-preserving machine learning in consumer electronics communication systems.

1) NIDS-FLGDP [36]: Network intrusion detection algorithm based on Gaussian differential privacy federated learning.
2) AADP [37]: An adaptive adjusted differential privacy federated learning method.
3) DP-DLCF [38]: A differential privacy-preserving deep learning caching framework for heterogeneous communication networks.
4) SSL-FL [39]: Self-sovereign identity-based privacy-preserving federated learning.
5) GeoPM-DMEIRL [40]: A deep inverse reinforcement learning security trajectory generation framework with serverless computing.
6) DPDA [41]: A novel differentially private domain adaptation framework called DPDA to achieve domain adaptation with consumer electronics assurance.
7) BCSEEDM [42]: A blockchain-based credible and secure educational consumer electronic data management scheme supporting for searchable encryption.

Fig. 2 shows the relationship between privacy loss accumulation and training epoch progression in educational deep learning systems, demonstrating how different privacy accounting mechanisms perform under extended training scenarios typical of next generation education consumer electronics communication systems. The analysis reveals the superior efficiency of zero-concentration differential privacy accounting compared to the traditional moment accountant. Standard differential privacy approaches when $\sigma_{\text{edu,max}} = \sigma_{\text{base}} = \sigma_{\text{edu},2} = 5$ and $\delta_{\text{edu}} = 10^{-2}$.

Fig. 2 illustrates privacy loss accumulation across training epochs using three accounting methods. The zero-concentrated differential privacy approach (blue circles) shows the slowest growth, reaching approximately 4.5 after one hundred epochs. The moment the accountant (red squares) accumulates faster, approaching 13 at one hundred epochs.

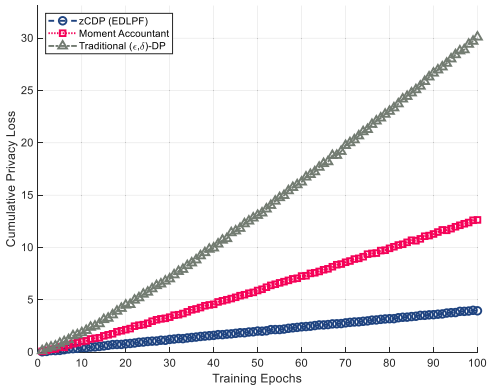

Fig. 2. Privacy loss vs. training epochs relationship.

Traditional differential privacy composition (gray triangles) grows linearly and most rapidly, exceeding 30 by one hundred epochs. This comparison demonstrates that zero-concentrated differential privacy enables substantially more training iterations under identical privacy budgets, permitting the extended training necessary for educational deep learning while maintaining rigorous mathematical guarantees against privacy violations.

B. Educational Framework Effectiveness Validation

Fig. 3 shows the validation of the EDLPF framework for effectiveness under different scale settings for the noise. It demonstrates that an adaptive differential privacy framework can uphold the utility of the consumer electronics model whilst simultaneously providing a meaningful privacy guarantee.

Fig. 3 validates ECEDL effectiveness across different noise scales using MNIST and Fashion-MNIST educational datasets. Each plot shows classification accuracy versus training epochs under three noise conditions: baseline without privacy, ECEDL with moderate noise, and ECEDL with higher noise. The MNIST results demonstrate accuracies exceeding 0.98 even with privacy protection. Fashion-MNIST shows similar patterns with slightly lower absolute accuracy due to increased task complexity. Both datasets reveal that ECEDL maintains stable convergence despite noise injection, with performance degradation remaining minimal until noise scales become substantial. The framework achieves practical utility-privacy trade-offs across diverse educational classification tasks.

The comparative analysis based on the performance between the EDLPF framework and the other six baselines from the MNIST training dataset, as given in Fig. 4, indicates that adaptive differential-private mechanisms for educational deep learning facilitated better convergence and sustained accuracy performances.

Fig. 4 compares ECEDL against seven baseline methods on MNIST educational data over one hundred training epochs. The ECEDL framework consistently achieves the highest accuracy, stabilizing above 0.99. BCSEEDM performs second-best, reaching approximately 0.97. DPDA and SSL-FL cluster around 0.96 accuracy. GeoPM-DMEIRL and DP-DLCF show moderate performance near 0.95. AADP and NIDS-FLGDP exhibit lower accuracy below 0.94. All methods demonstrate

Authorized licensed use limited to: Northeastern University. Downloaded on August 19,2026 at 00:41:31 UTC from IEEE Xplore. Restrictions apply.

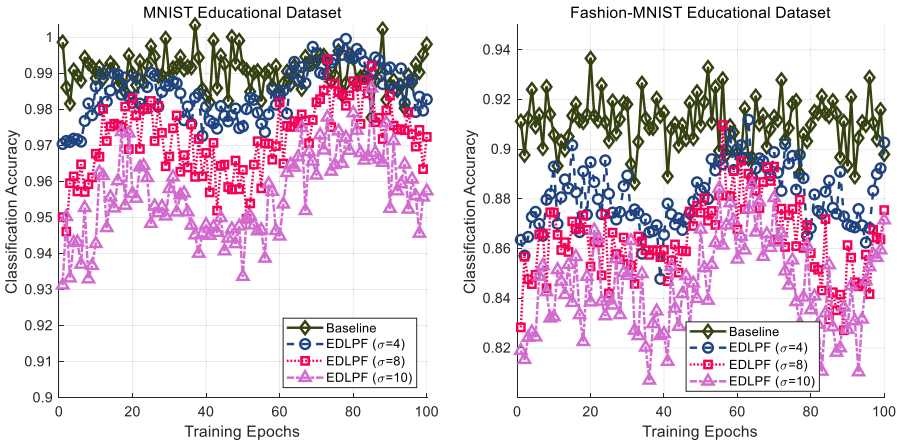

Fig. 3. Educational framework effectiveness validation.

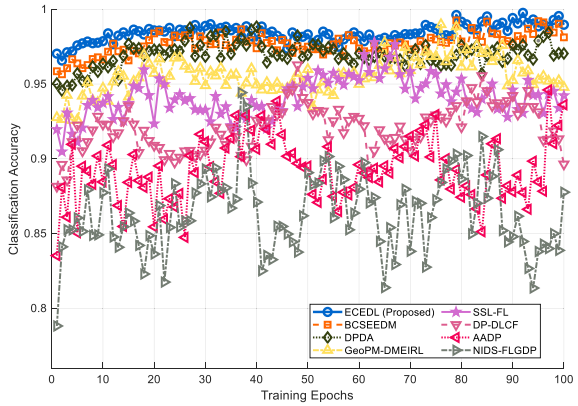

Fig. 4. MNIST consumer electronics educational dataset comparative performance.

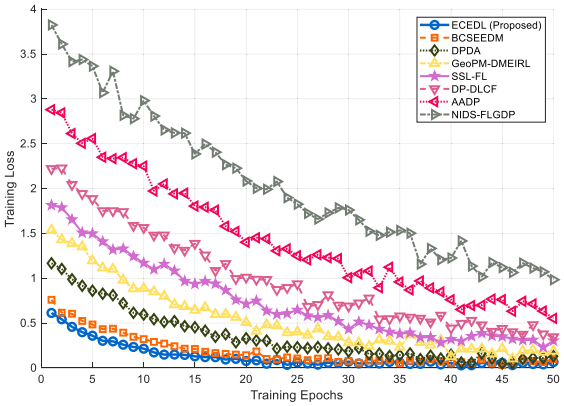

Fig. 6. Convergence rate analysis for consumer electronics educational systems.

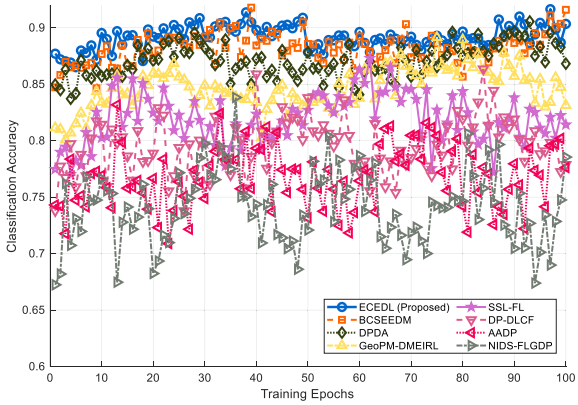

Fig. 5. Fashion-MNIST consumer electronics educational dataset comparative performance.

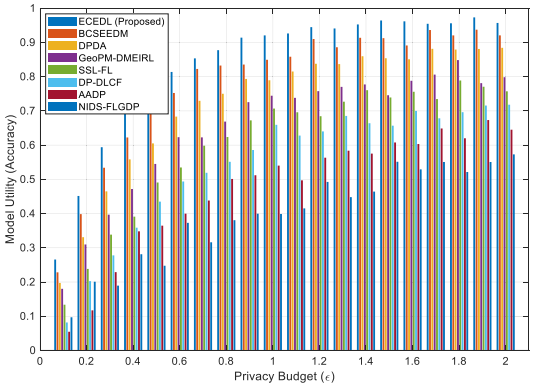

Fig. 7. Privacy budget efficiency analysis for consumer electronics educational systems.

convergence within forty epochs, though final accuracy levels differ substantially. The performance ordering remains stable throughout training, confirming ECEDL's superiority stems from fundamental design advantages rather than convergence speed differences alone.

A controlled performance characterization of Fashion-MNIST educational-context corpora, as in Fig. 5, corroborates the broad drought mitigation potentials of the EDLPF scheme. It is a somewhat unmistakable trajectory that dominates the incumbents of the canonical and adaptive might insinuate a broader operating envelope against synthetic stratifications of task-embedded patterns and synthetic learner behavior profiles. These distortions that functionally duplicate instability, often entrenched in emerging educational mediation networks, ensure that the EDLPF implementation is more durable and thus can offer a sustainable adoption even prior to life-cycle assistance from forthcoming architectures.

Authorized licensed use limited to: Northeastern University. Downloaded on August 19,2026 at 00:41:31 UTC from IEEE Xplore. Restrictions apply.

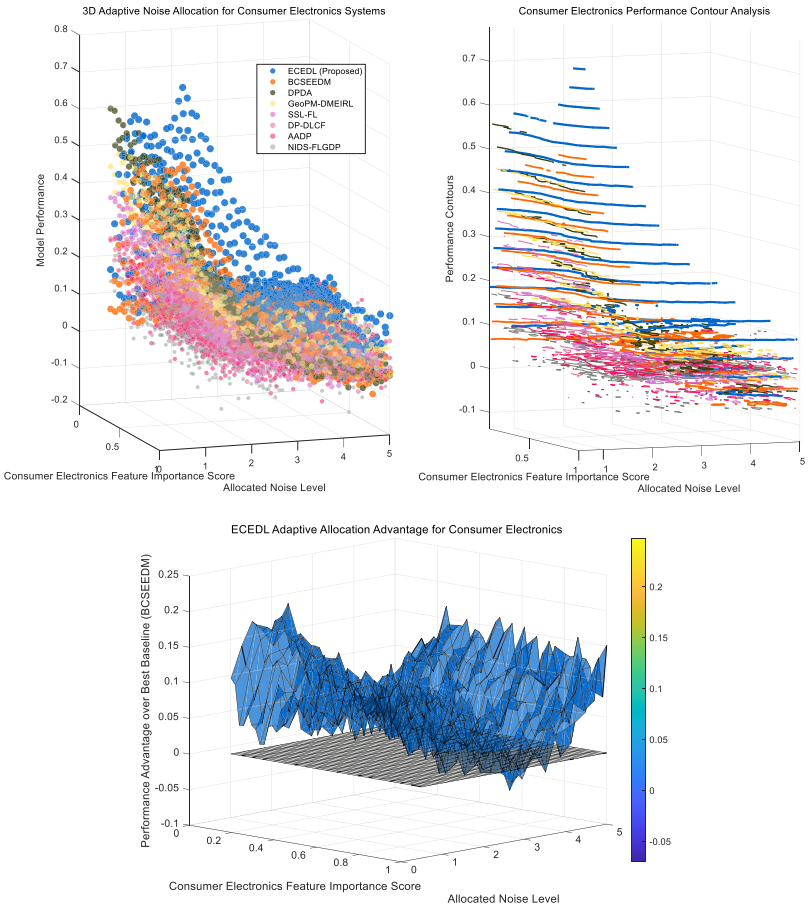

Fig. 8. Adaptive noise allocation effectiveness analysis for consumer electronics.

Benchmarking results in the Fashion-MNIST scheme have proven that EDLPF decisively outperforms reference algorithms by 2-4% in classification accuracy over consumer electronics educational datasets of consumer electronics complexity. Such sustained amplifications of accuracy beyond cohorts of increasing complexity intrinsic distributions narrow into the proof of the adaptive capability of the framework toward the plurality of teaching environments expected in modern educational communications systems. Besides that, further guarantees are provided by the monotonically flat convergence profile, ensuring stable performance evolutions in the consumer electronics setup since the EDLPF itself integrates trust and accountability models - a property that highly favorably corresponds to transparent communication on behalf of students and instructors.

Fig. 6 shows the convergence rate analysis of the EDLPF framework against baseline methods. It has been demonstrated that adaptive optimization algorithms tailored towards educational deep learning facilitate quicker and steadier convergence under privacy-preserving circumstances.

The convergence analysis shows a much quicker convergence achieved by the EDLPF framework compared to all other baseline methods, requiring around 15 epochs to reach a stabilized performance compared with 25-40 epochs required by the competing approaches. Such a level of convergence efficiency would directly act as a benefit to the data security modeling based on deep learning in next-generation education consumer electronics communication systems, wherein computational overhead could be reduced, and models could be updated frequently in such a dynamic educational environment in which student learning patterns change all the time.

Fig. 7 analyzes privacy budget efficiency, comparing how various frameworks utilize the allocated privacy budgets to achieve the best trade-offs between privacy protections and model utility in consumer electronics applications.

The analysis of differential privacy budgeting reveals that the EDLPF framework reaches the highest overall utility profile across all privacy levies, yielding an accuracy surplus of almost 15% compared to the most competitive baselines observing identical privacy guarantees. The aforementioned performance distinction thus acts as a strategic asset for educational entities in assuring higher privacy, but with observable better accuracy. This assures consumer electronics stakeholders, not to mention its critical importance with existing legislating climate and the higher priority of safeguarding consumer electronics-generated data.

Fig. 8 demonstrates the adaptive noise distribution efficiency study by visualizing the three-dimensional feature significance, noise levels of distribution, and model performance. This experiment shows that the smart noise distribution policy of EDLPF is superior to the uniform allocation of alternative strategies of the baseline methods.

Authorized licensed use limited to: Northeastern University. Downloaded on August 19,2026 at 00:41:31 UTC from IEEE Xplore. Restrictions apply.

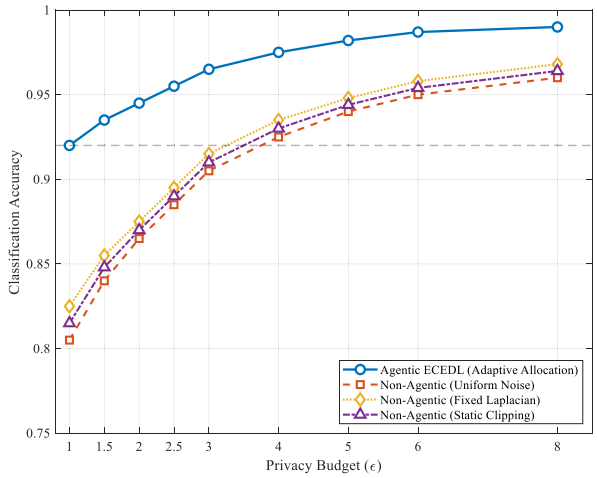

Fig. 9. Agentic vs non-agentic privacy systems performance.

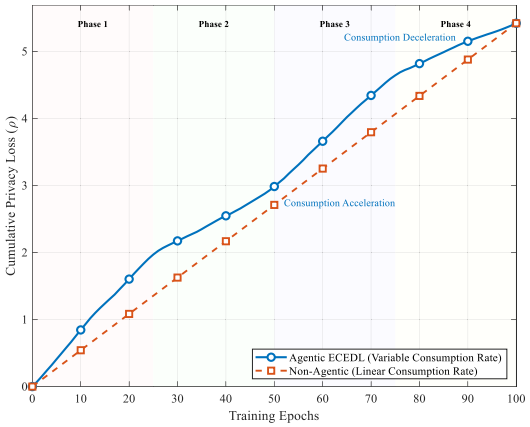

Fig. 10. Privacy budget consumption: Agentic vs non-agentic systems.

The adaptive noise allocation analysis shows that the intelligent feature-importance-based privacy distribution of EDLPF significantly improves uniform allocation strategies over the entire parameter space. This adaptive approach brings tremendous benefits to consumer electronics applications, where sensitive student characteristics get proportional protection and consumer electronics applications remain useful with their analysis capabilities.

C. Agentic AI Performance Analysis

Fig. 9 presents comparative performance between agentic and non-agentic privacy systems across increasing privacy budget constraints. The agentic ECEDL system maintains accuracy above ninety-two percent even at restrictive epsilon values of 2.0, demonstrating effective budget utilization through strategic allocation. Non-agentic systems applying uniform noise distributions show steeper accuracy degradation, dropping below eighty-five percent at equivalent privacy levels.

The mechanisms underlying agentic advantages become clearer through budget utilization analysis. Fig. 10 displays privacy budget consumption trajectories over training epochs, comparing agentic adaptive allocation against non-agentic fixed allocation. The agentic system exhibits variable consumption rates, accelerating budget use during training phases where sensitive features dominate gradients and decelerating when less-critical attributes drive learning.

Adaptation speed constitutes another dimension where agentic systems excel. We introduce a perturbation experiment where feature importance distributions shift abruptly at epoch fifty, simulating scenarios where educational models transition from demographic-focused early training to behavior-focused later training. The agentic ECEDL system detects this shift within five epochs through sensitivity monitoring and reallocates privacy budgets accordingly. Accuracy recovers to within two percent of optimal allocation by epoch sixty. Non-agentic systems cannot respond to importance shifts, maintaining outdated allocations that misallocate resources throughout remaining training. This results in sustained accuracy penalties averaging eight percent compared to agentic adaptation.

Stepping back from individual metrics, a coherent picture emerges. The ECEDL framework demonstrates that meaningful privacy protection for educational deep learning is achievable without abandoning the analytical power these systems provide. The adaptive mechanisms prove their worth most clearly in heterogeneous scenarios—precisely the conditions real educational platforms encounter. When student data sensitivity varies across populations, and learning contexts shift between subjects, static privacy approaches waste budget where none is needed while under-protecting where risks concentrate.

V. Conclusion

This study presented the ECEDL framework, addressing critical privacy vulnerabilities in next generation consumer electronics-powered education communication systems through adaptive differential privacy mechanisms specifically designed for heterogeneous educational technology environments. Experimental validation across diverse datasets demonstrated that the ECEDL framework consistently outperforms seven established baseline methods, achieving ten to fifteen percent improvements in classification accuracy while maintaining equivalent or superior privacy guarantees. However, real-world consumer electronics educational data typically exhibits temporal dependencies reflecting learning progression, hierarchical relationships between educational concepts, and multi-modal characteristics combining interaction logs, assessment data, behavioral patterns, and potentially biometric information from wearable devices. The integration of quantum-resistant privacy mechanisms becomes increasingly important as quantum computing advances threaten current cryptographic foundations underlying consumer electronics security architectures.

Several research directions emerge from this work. Theoretically, alternative sensitivity measures beyond SHAP-based attribution warrant exploration, including information-theoretic approaches that capture feature dependencies more completely. Tighter composition theorems for adaptive mechanisms could recover privacy budget currently sacrificed to conservative worst-case assumptions. Engineering challenges include developing streaming privacy accounting for continuous student data arrival, computational optimizations for institutional-scale deployment processing millions of records,

Authorized licensed use limited to: Northeastern University. Downloaded on August 19,2026 at 00:41:31 UTC from IEEE Xplore. Restrictions apply.

and federated learning integration enabling cross-institutional analytics while maintaining data sovereignty. Application expansions should address recommendation systems handling implicit feedback, natural language processing on student-generated text, and longitudinal trajectory analysis accounting for repeated observations.

REFERENCES

[1] J. Liu, Z. Guan, X. Wang, and Z. Yang, "Fusing interaction transition features and cross-attention for knowledge tracing in E-learning systems," *IEEE Trans. Consum. Electron.*, vol. 71, no. 1, pp. 761–772, Feb. 2025.

[2] F. A. M. Al-Yarimi, R. Salah, and K. Mohamoud, "Blockchain-driven secure data sharing framework for edge computing networks," *Tsinghua Sci. Technol.*, vol. 30, no. 3, pp. 978–997, Jun. 2025.

[3] R. Ma and X. Chen, "Intelligent education evaluation mechanism on ideology and politics with 5G: PSO-driven edge computing approach," *Wireless Netw.*, vol. 29, no. 2, pp. 685–696, Feb. 2022.

[4] G. Lee et al., "Multimodality of AI for education: Toward artificial general intelligence," *IEEE Trans. Learn. Technol.*, vol. 18, pp. 666–683, 2025.

[5] R. Cerezo, J.-A. Lara, R. Azevedo, and C. Romero, "Reviewing the differences between learning analytics and educational data mining: Towards educational data science," *Comput. Hum. Behav.*, vol. 154, May 2024, Art. no. 108155.

[6] A. Kumi-Yeboah, Y. Kim, B. Yankson, S. Aikins, and Y. A. Dadson, "Diverse students' perspectives on privacy and technology integration in higher education," *Brit. J. Educ. Technol.*, vol. 54, no. 6, pp. 1671–1692, Aug. 2023.

[7] P. Petrov, I. Kuyumdzhiev, R. Malkawi, G. Dimitrov, and J. Jordanov, "Digitalization of educational services with regard to policy for information security," *TEM J.*, vol. 11, pp. 1093–1102, Aug. 2022.

[8] J. Chen, C. Hu, W. Sheng, R. Li, R. Zhao, and J. Yu, "A trust-based personalized differential privacy guarantees for online social networks," *IEEE Trans. Netw. Service Manage.*, vol. 22, no. 4, pp. 3213–3227, Aug. 2025.

[9] A. P. Muniyandi et al., "Privacy preserved reinforcement learning model using generative AI for personalized E-learning," *IEEE Trans. Consum. Electron.*, vol. 70, no. 3, pp. 6157–6165, Aug. 2024.

[10] Y. Pei and G. Lu, "Design of an intelligent educational evaluation system using deep learning," *IEEE Access*, vol. 11, pp. 29790–29799, 2023.

[11] K. Tajiri and R. Kawahara, "Optimization of data and model transfer for federated learning to manage large-scale network," *IEEE Trans. Netw. Service Manage.*, vol. 22, no. 2, pp. 958–973, Apr. 2025.

[12] P. N. Bathula and M. Sreenivasulu, "An integrated blockchain framework for secure data sharing in IoT fog computing," *Tsinghua Sci. Technol.*, vol. 30, no. 3, pp. 957–977, Jun. 2025.

[13] A. Al-Motrif, "Enhancing learning experiences for students with learning disabilities through digital pedagogies: Insights from Saudi schools," *Interact. Learn. Environments*, vol. 33, no. 5, pp. 3635–3662, Jan. 2025.

[14] M. Uddin, M. Abdelhaq, H. Alsufyani, F. Alanazi, and R. Alsaqour, "SEIMR: Secure and ethical IoT mixed reality integration framework for metaverse experiences," *IEEE Trans. Consum. Electron.*, vol. 71, no. 2, pp. 5691–5698, May 2025.

[15] X. Liu, "Blockchain-enabled collaborative edge computing for intelligent education systems using social IoT," *Int. J. Distrib. Syst. Technol.*, vol. 13, no. 7, pp. 1–19, Sep. 2022.

[16] X. Li, Z. Zhou, Q. He, Z. Shi, W. Gaaloul, and S. Yangui, "Re-scheduling IoT services in edge networks," *IEEE Trans. Netw. Service Manage.*, vol. 20, no. 3, pp. 3233–3246, Sep. 2023.

[17] J. Lv, B.-G. Kim, B. D. Parameshachari, A. Slowik, and K. Li, "Large model-driven hyperscale healthcare data fusion analysis in complex multi-sensors," *Inf. Fusion*, vol. 115, Mar. 2025, Art. no. 102780.

[18] M. Nickl et al., "Effects of real-time adaptivity of scaffolding: Supporting pre-service mathematics teachers' assessment skills in simulations," *Learn. Instruct.*, vol. 94, Dec. 2024, Art. no. 101994.

[19] J. Sun et al., "Weighted heterogeneous graph-based three-view contrastive learning for knowledge tracing in personalized e-learning systems," *IEEE Trans. Consum. Electron.*, vol. 70, no. 1, pp. 2838–2847, Jul. 2024.

[20] C. Wangang, "Design and implementation of an educational information management system using deep learning and wireless communication," *Mobile Netw. Appl.*, vol. 28, no. 6, pp. 2138–2148, Dec. 2023.

[21] P. Huang, G. Liao, and J. Ren, "The application of AES-SM2 hybrid encryption algorithm in big data security and privacy protection," *Int. J. Adv. Comput. Sci. Appl.*, vol. 15, no. 6, pp. 989–997, Jun. 2024.

[22] Q. Liu, R. Shakya, J. Jovanovic, M. Khalil, and J. de la Hoz-Ruiz, "Ensuring privacy through synthetic data generation in education," *Brit. J. Educ. Technol.*, vol. 56, no. 3, pp. 1053–1073, May 2025.

[23] J. Soria-Comas and J. Domingo-Ferrer, "Optimal data-independent noise for differential privacy," *Inf. Sci.*, vol. 250, pp. 200–214, Nov. 2013.

[24] L. T. Phong, Y. Aono, T. Hayashi, L. Wang, and S. Moriai, "Privacy-preserving deep learning via additively homomorphic encryption," *IEEE Trans. Inf. Forensics Security*, vol. 13, no. 5, pp. 1333–1345, May 2018.

[25] J. Zhang, G. Cormode, C. M. Procopiuc, D. Srivastava, and X. Xiao, "PrivBayes: Private data release via Bayesian networks," *ACM Trans. Database Syst.*, vol. 42, no. 4, pp. 1–41, Nov. 2017.

[26] M. E. Gürsoy, A. Inan, M. E. Nergiz, and Y. Saygın, "Differentially private nearest neighbor classification," *Data Mining Knowl. Discovery*, vol. 31, no. 5, pp. 1544–1575, Sep. 2017.

[27] Z. Yan, G. Li, and J. Liu, "Private rank aggregation under local differential privacy," *Int. J. Intell. Syst.*, vol. 35, no. 10, pp. 1492–1519, Oct. 2020.

[28] Y. Zhao, J. T. Du, and J. Chen, "Scenario-based adaptations of differential privacy: A technical survey," *ACM Comput. Surv.*, vol. 56, no. 8, pp. 1–39, Aug. 2024.

[29] M. Chakraborty, W. Pal, S. Bandyopadhyay, and U. Maulik, "A survey on multi-objective based parameter optimization for deep learning," *Comput. Sci.*, vol. 24, no. 3, pp. 321–353, Oct. 2023.

[30] X. Liu, J. Yang, S. Chen, H. He, and X. Jiang, "PRFL: Achieving efficient robust aggregation in privacy-preserving federated learning," *IEEE Trans. Netw. Serv. Manag.*, vol. 22, no. 2, pp. 2138–2155, Apr. 2024.

[31] K. Lin et al., "TAG: Teacher-advice mechanism with Gaussian process for reinforcement learning," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 9, pp. 12419–12433, Sep. 2024.

[32] X. Wang, H. Zhang, M. Yang, X. Wu, and P. Cheng, "Privacy-preserving collaborative learning: A scheme providing heterogeneous protection," *IEEE Internet Things J.*, vol. 11, no. 2, pp. 1840–1853, Jan. 2024.

[33] X. Hu, M. Zhu, Z. Feng, and L. Stanković, "Manifold-based Shapley explanations for high dimensional correlated features," *Neural Netw.*, vol. 180, Dec. 2024, Art. no. 106634.

[34] L. G. McKnight, C. Jaiswal, I. W. AlHmoud, and B. Gokaraju, "A dataset of deep learning performance from cross-base data encoding on MNIST and MNIST-C," *Data Brief*, vol. 57, Dec. 2024, Art. no. 111194.

[35] R. I. Mukhamediev, "State-of-the-art results with the fashion-MNIST dataset," *Mathematics*, vol. 12, no. 20, p. 3174, Oct. 2024.

[36] J. Du and K. Yang, "NIDS-FLGDP: Network intrusion detection algorithm based on Gaussian differential privacy federated learning," *J. Circuits, Syst. Comput.*, vol. 33, no. 3, Feb. 2024, Art. no. 2450048.

[37] Y. Cheng, W. Li, S. Qin, and T. Tu, "Differential privacy federated learning based on adaptive adjustment," *Comput., Mater. Continua*, vol. 82, no. 3, pp. 4777–4795, Mar. 2025.

[38] H. Wang, X. Zhang, Y. Xia, and X. Wu, "A differential privacy-preserving deep learning caching framework for heterogeneous communication network systems," *Int. J. Intell. Syst.*, vol. 37, no. 12, pp. 11142–11166, Dec. 2022.

[39] R. Ul Haque, A. S. M. T. Hasan, and M. A. M. Al-Hababi, "SSL-FL: Self-sovereign identity based privacy-preserving federated learning," *J. Parallel Distrib. Comput.*, vol. 191, Sep. 2024, Art. no. n104907.

[40] Y.-R. Huang, J. Zhang, H.-M. Hou, X.-C. Ye, and Y. Chen, "GeoPM-DMEIRL: A deep inverse reinforcement learning security trajectory generation framework with serverless computing," *Future Gener. Comput. Syst.*, vol. 154, pp. 123–139, May 2024.

[41] Q. Wang, Z. Li, Q. Zou, L. Zhao, and S. Wang, "Deep domain adaptation with differential privacy," *IEEE Trans. Inf. Forensics Security*, vol. 15, pp. 3093–3106, 2020.

[42] Z. Li and Z. Ma, "A blockchain-based credible and secure education experience data management scheme supporting for searchable encryption," *China Commun.*, vol. 18, no. 6, pp. 172–183, Jun. 2021.

Authorized licensed use limited to: Northeastern University. Downloaded on August 19,2026 at 00:41:31 UTC from IEEE Xplore. Restrictions apply.