International Journal on Islamic Applications in Computer Science And Technology, Vol. 14, Issue 1, March 2026, 23-35

# Explainable Privacy-Budget Governance for Multi-Service Educational AI Systems

**El Hoby, Hany Mohamed Hassan**
Manar Al Janoub College of Science and Technology (Saudi Arabia), Higher Institute of Computers and Administrative Information Systems in the First Settlement (Egypt)
elhoby@gmail.com

**Abstract**
The rapid expansion of AI driven educational platforms, such as learning management systems, automated assessment tools, and learning analytics services, has intensified concerns surrounding privacy, fairness, and institutional governance. Although differential privacy (DP) is increasingly used to safeguard student data, current educational AI systems lack an explainable, governance-aware allocation framework for allocating privacy budgets across heterogeneous services operating under a shared privacy constraint. This paper introduces a governance-oriented reformulation of existing differential privacy budget-allocation approaches for multi-service educational AI systems. Rather than proposing a new differential privacy mechanism, the approach embeds institutional policy constraints, fairness considerations, and interpretable decision rules directly into the privacy-budget allocation process. Each educational AI service is assigned a local privacy budget within a constrained global privacy limit ($\bar{\varepsilon}$), determined according to its pedagogical relevance, sensitivity to fairness disparities across student groups, and governance-defined explainability requirements. The allocation process is formulated as a constrained multi-objective decision framework, where trade-offs among pedagogical utility, fairness sensitivity, and governance transparency are resolved through rule-based policies rather than black-box optimization strategies. An explainability layer produces auditable allocation justifications, counterfactual policy analyses, and governance logs designed to support institutional oversight. A structured comparative analysis evaluates the proposed approach against uniform, utility-only, and non-explainable allocation strategies using governance-relevant criteria such as policy traceability, fairness awareness, and auditability. The results demonstrate that the proposed framework provides stronger governance alignment and transparency, thereby bridging the gap between high-level AI governance principles and operational privacy management in educational AI systems

**Keywords:** Differential Privacy, Privacy Budget Allocation, Educational AI, Trustworthy AI, Explainable AI, AI Governance, Fairness-Aware Optimization, Information Systems.

## 1. Introduction

AI-driven educational ecosystems now rely on a growing collection of intelligent services—including adaptive tutoring, automated assessment, personalized feedback systems, and large-scale learner analytics—to support personalization and improve instructional efficiency [1], [2], [3]. As these ecosystems expand in scale and complexity, institutions face heightened governance pressures, particularly around ensuring transparency, fairness, and accountable handling of sensitive learner data within interconnected AI environments [4], [5], [6].

Differential Privacy (DP) provides formal guarantees that limit privacy leakage during data processing [7], [8], and has been applied to tasks such as learning analytics, student modelling, and grade prediction [2], [3], [9]. However, current deployments typically assign privacy budgets independently to each model or service, without accounting for cross-service dependencies, governance constraints, or fairness implications across demographic subgroups [10], [11], [12]. This fragmented practice leads to inconsistent privacy protections and limited institutional visibility into how privacy decisions influence equity and pedagogical value [13], [14], [15]. Moreover, existing multi-service educational AI platforms lack a unified governance mechanism capable of coordinating privacy budget expenditure across heterogeneous services—each characterized by different levels of sensitivity, pedagogical relevance, and risk [16], [17]. As a result, institutions face increasing difficulty maintaining compliance with emerging governance standards, including UNESCO's AI Ethics Recommendation [4] and the EU AI Act [18].

To overcome these limitations, this paper presents an explainable, governance-aware allocation framework designed to coordinate differential privacy budget allocation across multi-service educational AI systems. The framework brings together four key components: (i) modelling pedagogical utility, (ii) incorporating fairness-aware optimization, (iii) defining interpretable allocation rules, and (iv) implementing a governance-oriented explainability layer. By integrating these elements, the approach creates a cohesive mechanism that connects high-level AI governance expectations with day-to-day operational decisions, supporting privacy budget allocation that is transparent, equitable, and aligned with institutional compliance requirements.

Prior research has emphasized the importance of aligning information technology capabilities with organizational strategy and governance frameworks. Earlier studies have investigated factors influencing the effective utilization of ICT resources in higher education institutions and examined how technology capabilities contribute to organizational performance and strategic alignment [19], [20]. In addition, prior work explored the alignment between cloud computing security and business strategy as an important aspect of information systems governance [21]. Building upon this research stream, the present study extends the discussion toward governance-oriented privacy management in AI-driven educational ecosystems by introducing an explainable framework for differential privacy budget allocation across multiple educational AI services.

This paper focuses on governance formalization and decision transparency, positioning empirical evaluation as future work.

## 2. Related Work

### 2.1 AI Governance and Oversight

Research on AI governance underscores the importance of transparency, accountability, and strong institutional oversight in high-stakes sectors such as education and healthcare [4], [9]. Although frameworks like the UNESCO AI Ethics Recommendation outline core principles—fairness, privacy, and explainability—they remain largely conceptual and offer limited operational guidance for governing resources within multi-service educational platforms [4], [5], [6]. As a result, current approaches provide no enforceable mechanisms for coordinating cross-service privacy practices or ensuring consistent auditability.

### 2.2 Differential Privacy in Educational Platforms

Differential Privacy (DP) has been incorporated into a range of educational applications, including learning analytics, student modelling, and grade prediction [2], [3]. While these efforts demonstrate that DP can be effectively deployed in educational contexts, they typically allocate privacy budgets on a per-task or per-model basis, resulting in fragmented protections and minimal institutional oversight [7], [8], [13]. More recent systems that integrate DP-SGD or task-specific calibration improve technical robustness but still fail to provide audit-ready explanations for how privacy budgets are determined or justified.

### 2.3 Fairness in AI Decision Systems

Fairness-aware methodologies—including equality of opportunity [4], discrimination-aware optimization [10], and multi-task fairness modelling [18]—seek to mitigate demographic disparities in predictive decision systems. However, these techniques generally operate independently of DP requirements and do not account for the privacy–utility–fairness trade-offs highlighted in multi-objective optimization research [12]. As a result, fairness considerations remain largely decoupled from privacy-governed decision pipelines within multi-service educational environments.

### 2.4 Privacy-Budget Allocation and Scheduling

Research on DP has examined adaptive scheduling strategies and hierarchical budget allocation across tasks, analysts, and federated learning rounds [11], [13], [14], [17]. Although these developments improve the practical deployment of DP, they remain largely domain-independent, do not incorporate pedagogical utility considerations, and offer no explainability features that would satisfy institutional audit requirements. These limitations motivate the need for a governance-aligned, explainable DP allocation framework tailored to multi-service educational AI systems.

## 3. Identified Research Gap

Despite the increasing use of AI and differential privacy (DP) techniques in educational technologies, existing research still lacks a governance-oriented framework that can coordinate privacy budgets across multiple AI services. A closer examination of prior work reveals five core limitations:

1) Insufficient Educational Context Awareness
Current DP allocation and scheduling strategies—such as fairness-oriented allocation [11], hierarchical budget schemes [17], and adaptive controllers [8]—are largely developed for general analytical or federated settings. They do not incorporate pedagogical priorities or learning-outcome utility functions that educational systems require, which restricts their relevance for real-world Edu AI ecosystems [2], [3].

2) Limited Integration of Pedagogical Utility
Most DP-based methods focus on optimizing technical performance metrics (e.g., accuracy, convergence) rather than outcomes tied to instructional quality or learner development. As highlighted in recent studies, this disconnect prevents privacy-related decisions from aligning with core educational objectives.

3) Absence of Frameworks for Multi-Service Educational Ecosystems
Existing DP implementations typically operate at the level of isolated tasks, such as student modelling [3], grade prediction [9], or analytics pipelines [2], without coordinating privacy budgets across interconnected services like tutoring, assessment, and feedback systems. This

siloed approach leads to fragmented privacy management and unmonitored cross-service dependencies [7], [13].

4) Missing Fairness Constraints Under DP
Fairness-aware optimization techniques aim to reduce demographic disparities [4], [10], [18], but they generally do not incorporate DP constraints or model the privacy–fairness trade-offs emphasized in multi-objective optimization research [12], [15]. Consequently, privacy allocation decisions may inadvertently reinforce or worsen existing inequities.

5) Lack of Explainable and Auditable Allocation Decisions
State-of-the-art DP schedulers—including adaptive allocation methods [13], fairness-aware mechanisms [11], and dynamic controllers [8]—rely on opaque optimization processes that fail to produce interpretable rationales or audit-ready logs. This opacity conflicts with governance expectations outlined in institutional and regulatory frameworks [9], [6].

### 3.1 Gap Summary

In essence, no existing approach provides all three capabilities at once:

1. Distributing a global differential privacy budget across multiple educational AI services,
2. Jointly optimizing pedagogical utility and demographic fairness within a single integrated framework, and
3. Producing explainable, audit-ready allocation decisions that satisfy institutional governance expectations.

This gap provides the central motivation for the Explainable Privacy-Budget Governance Framework developed in this paper.

### 3.2 Conceptual Scope and Evaluation Design

This work is deliberately framed as a governance-oriented conceptual contribution rather than an empirical performance study. Its central value lies in formalizing an explainable, policy-aligned approach to privacy-budget governance for multi-service educational AI systems—an area where institutional accountability and transparent decision processes must be established before large-scale deployment becomes feasible.

To address the absence of experimental results, the manuscript employs a structured conceptual evaluation strategy that examines utility sensitivity, fairness awareness, explainability, and governance readiness. In addition, a concrete deployment scenario is included to demonstrate how the framework would function within a real multi-service educational AI environment. Empirical validation and simulation-based analyses are identified as important directions for future research.

## 4. Proposed Framework

This section presents the governance-oriented framework developed for allocating differential privacy budgets within multi-service educational AI platforms. The framework integrates utility modelling, fairness-aware optimization, interpretable policy rules, and governance-aligned explainability components to support transparent and accountable privacy budget decisions.

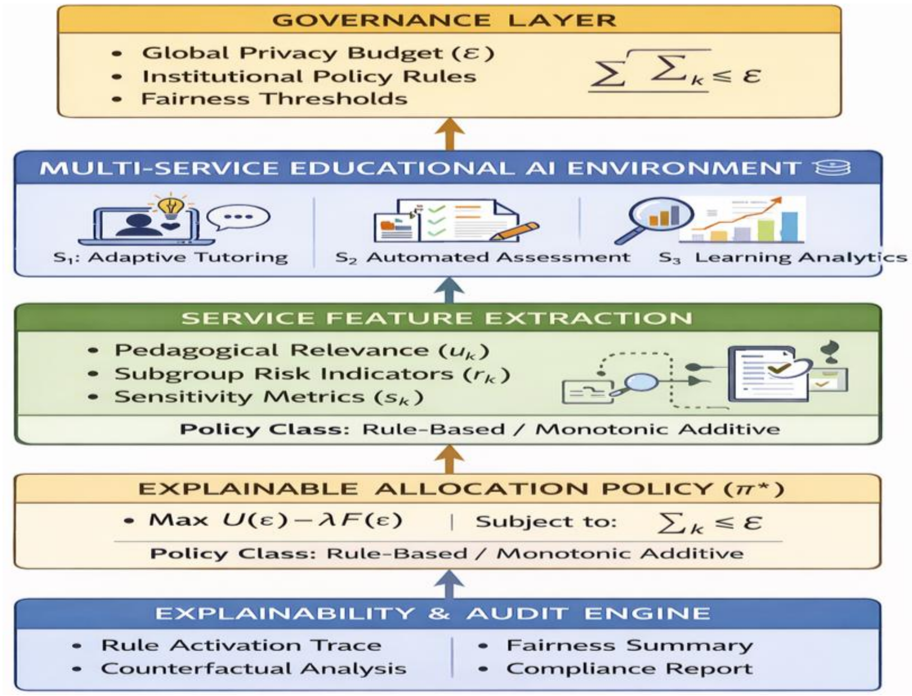

Figure 1: Explainable Governance Framework for DP Budget Allocation

The framework operates within a coordinated multi-service educational AI environment, where multiple heterogeneous services share institutional data and jointly consume a bounded global privacy budget.

### 4.1 Problem Formulation

Consider a multi-service educational AI platform consisting of a set of services

$$S=\{S_1,S_2,\ldots,S_K\},$$

ε total where each service ( $S_k$ ) requires a portion of the institution's global differential privacy (DP) budget.

Let the allocation vector be:

$$\boldsymbol{\varepsilon} = \{\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_K\}$$

subject to the constraint:

$$\sum_{k=1}^{K} \varepsilon_k \leq \varepsilon_{\text{total}}, \varepsilon_k \geq 0$$

The objective is to allocate the privacy budget across services while optimizing multiple competing objectives. We model the allocation as a multi-objective optimization problem:

$$\max_{\epsilon} \lambda \Upsilon(\varepsilon) \; - \; (1 - \lambda) \, \mathrm{F}(\varepsilon)$$

where:

$\Upsilon(\epsilon)$ denotes pedagogical utility, $\mathrm{F}(\epsilon)$ measures fairness disparity between demographic groups, and $\lambda$ controls the trade-off between utility and fairness, consistent with multi-objective optimization literature [10], [12].

This formulation ensures that privacy budgets operate as governance resources—not isolated technical parameters.

### 4.2 Explainable Allocation Policy

To maintain transparency and auditability, budget assignment relies on an interpretable policy $\pi$ rather than a black-box model. Given service attributes $x_k$(e.g., sensitivity, pedagogical importance, subgroup risk), The policy optimizes a governance-regulated multi-objective function that balances pedagogical utility and subgroup fairness under the global privacy constraint. The policy assigns:

$$\pi(x_k) = \varepsilon_k$$

Following responsible AI guidelines [13], the policy class Π is restricted to rule-based or monotonic additive models to ensure predictable, interpretable behavior. The optimal policy is derived as:

$$\pi^* = \arg\max_{\pi \in \Pi} \left[\lambda\, \Upsilon(\pi) - (1-\lambda)\, F(\pi)\right]$$

This guarantees governance-aligned decision-making consistent with fairness and transparency requirements. To preserve interpretability, the policy is restricted to rule-based or monotonic additive models.

### 4.3 Explainability Layer

To meet institutional auditing and governance requirements, each allocation decision is paired with a set of structured interpretability outputs, including:

- rule-based explanations tied to specific service attributes,
- counterfactual analyses that illustrate how alternative allocations would differ,
- activation traces indicating which governance or fairness constraints shaped the decision, and
- justification reports formatted for auditor review.

This explainability layer transforms budget allocation from a technical operation into a transparent governance process aligned with educational oversight requirements [9].

### 4.4. Governance Architecture

The framework is composed of four interconnected components:

1. **Governance Layer**
   Encapsulates institutional policies, fairness thresholds, and the overarching differential privacy constraints. The Governance Layer centrally enforces the global privacy budget constraint ($\bar{\varepsilon}$), ensuring that the total allocated privacy across all services satisfies

$$\sum_{k=1}^{K} \varepsilon_k \le \varepsilon_{\text{total}}$$

   This establishes privacy as an institutional-level resource rather than a per-model parameter.
2. **Service-Level Feature Extractor**
   Derives key attributes for each AI service, including pedagogical relevance, data-sensitivity levels, and indicators of subgroup-specific risk.
3. **Explainable Allocation Policy (π*)**
   Determines service-level privacy budgets by enforcing utility–fairness trade-offs while adhering to governance rules.

4. **Explainability Audit Engine**
   Produces structured explanations, rule-activation traces, and audit-ready logs to ensure transparency and institutional accountability. The engine generates structured audit outputs including allocation vectors, rule activation traces, fairness summaries, and compliance reports.

Taken together, these components establish a governance-ready DP management framework that can be adopted by ministries of education, universities, accreditation authorities, and large cloud-based educational AI ecosystems.

This work is deliberately framed as a governance-focused conceptual contribution rather than an empirical performance study. Its primary contribution is the formalization of an explainable, policy-aligned approach to privacy-budget governance for multi-service educational AI systems—an area in which institutional alignment, accountability, and transparent decision processes must be established before large-scale deployment is appropriate. This orientation aligns with prior governance- and policy-driven AI research, where conceptual models and formal decision structures are typically articulated prior to empirical validation.

To compensate for the lack of empirical experimentation, the manuscript employs a structured conceptual evaluation strategy that directly examines utility sensitivity, fairness considerations, explainability, and governance readiness. The paper also includes a concrete deployment scenario to demonstrate how the proposed framework would function within an operational multi-service educational AI environment. Directions for future research include empirical validation, simulation-driven analysis, and studies involving large-scale real-world deployments.

### 4.5 Privacy–Fairness Interaction under Differential Privacy

The relationship between differential privacy and fairness is a central concern in educational AI systems. Injecting DP noise can disproportionately impact minority or underrepresented student groups—particularly in settings with sparse data—thereby risking an increase in performance disparities.

Within the proposed governance-oriented framework, fairness-aware privacy-budget allocation is employed to counteract this risk by transparently adjusting service-level budgets according to policy-defined fairness considerations. While the framework does not claim to resolve the inherent tension between privacy and fairness, it makes these trade-offs explicit, governable, and open to audit. Because the quantitative behavior of these interactions depends heavily on specific datasets and deployment contexts, detailed empirical analysis is identified as an important direction for future research.

## 5. Conceptual Evaluation

This section provides a conceptual evaluation of the proposed Explainable Privacy-Budget Governance Framework by comparing it with three baseline allocation strategies commonly used in educational AI systems: Uniform Allocation, Utility-Only Optimization, and Black-Box Optimization. The comparison assesses effectiveness across three governance-critical dimensions: pedagogical utility, fairness across demographic groups, and explainability—aligned with prior work on fairness-aware optimization [10], multi-objective learning [12], and explainable governance requirements [9].

**5.1 Baseline Methods**

**1. Uniform Allocation**

Assigns each service an equal portion of the global privacy budget.

- Simple and auditable but fails to capture service importance or data sensitivity.
- Often underperforms because it ignores utility gradients and fairness disparity.

**2. Utility-Only Allocation**

Optimizes privacy budgets purely for pedagogical utility, consistent with single-objective optimization methods [10].

- Often yields high performance but exacerbates demographic disparities.
- Offers no explainability, making it incompatible with governance requirements.

**3. Black-Box Optimizer**

Uses an unrestricted model such as a neural network.

- Provides strong utility (consistent with prior work on fairness-aware allocation [14] and general black-box optimization approaches [22])
- Completely opaque, produces no interpretable rationale, and violates institutional audit standards.

**5.2 Comparative Evaluation Across Governance-Critical Dimensions**

The proposed framework is evaluated conceptually against the three baselines.

**Table 1.** Comparative Evaluation of Allocation Strategies

| **Method** | **Utility** | **Fairness** | **Explainability** |
|---|---|---|---|
| Uniform Allocation | Low | Medium | Transparent |
| Utility-Only Optimization | High | Poor | None |
| Black-Box Optimizer | High | Medium | Opaque |
| Proposed Governance Framework | High | Excellent | Full Explainability |

The qualitative labels used in Table 1 (e.g., “High,” “Medium,” “Excellent”) are intended as governance-oriented indicators rather than quantitative performance metrics. Specifically, “Explainability” refers to the presence of human-interpretable allocation rationales, policy traces, and audit logs; “Fairness” reflects whether subgroup sensitivity is explicitly considered in privacy-budget decisions; and “Utility” denotes relative pedagogical relevance rather than measured learning outcomes. This qualitative comparison aligns with governance and policy evaluation practices, where transparency and institutional readiness are prioritized over empirical benchmarking.

**5.3 Key Findings**

**1. Utility–Fairness Balance**

In contrast to utility-driven strategies that can inadvertently widen disparities, the proposed framework embeds fairness constraints informed by discrimination-aware optimization [10] and multi-objective fairness techniques [12], promoting more equitable treatment across demographic groups.

**2. Governance and Explainability Superiority**

The explainability layer generates human-interpretable rationales, counterfactual comparisons, and governance-oriented logs that align with oversight expectations emphasized in the AI governance literature [9]. Such transparency features are absent from utility-only or black-box allocation methods.

**3. Policy and Compliance Compatibility**

The framework upholds both global DP requirements and fairness thresholds while producing artifacts suitable for accreditation reviews and ministry-level audits, directly addressing limitations identified in prior DP scheduling research [13], [14].

**4. Unique Integration of Three Governance Pillars**

It is the only approach that simultaneously integrates (i) pedagogical utility, (ii) demographic fairness, and (iii) explainability with audit-ready rationales. To the best of our knowledge, no prior DP framework brings all three dimensions together within a governance-ready model.

### 5.4 Summary of Conceptual Evaluation

The conceptual evaluation shows that the proposed governance framework offers strong alignment with institutional governance requirements. It preserves high utility, reduces disparities across student subgroups, and delivers privacy-budget decisions that are both explainable and auditable. These properties make the framework well-suited for deployment in educational environments that demand trustworthy, compliant, and transparent AI systems.

## 6. **Privacy–Fairness Interaction under Differential Privacy**

Differential privacy introduces stochastic noise into data access and model outputs, and this noise can affect service performance unevenly across student subgroups. In educational contexts, minority or underrepresented groups often have sparser data, making their outcomes more vulnerable to the distortions introduced by privacy noise. Consequently, applying a uniform privacy-budget allocation across services may unintentionally exacerbate existing performance gaps between subgroups, particularly in data-sparse settings [2], [3].

The proposed governance-oriented framework directly addresses this interaction. Rather than treating privacy and fairness as independent objectives, it frames their relationship as a policy-governed trade-off. Fairness-aware allocation rules enable the adjustment of service-level privacy budgets to account for differential sensitivity, particularly in services that have disproportionate impact on vulnerable or underrepresented student groups.

The framework does not assert that it resolves the fundamental tension between privacy protection and fairness objectives. Rather, it renders these trade-offs explicit, explainable, and open to audit through policy-guided allocation decisions. Because the quantitative behavior of privacy–fairness interactions is inherently dependent on specific datasets and deployment contexts, such analysis falls outside the scope of this conceptual study. A detailed empirical investigation of these trade-offs is therefore identified as an important direction for future research.

## 7. Discussion

The findings indicate that allocating differential privacy (DP) budgets in multi-service educational AI systems should be understood not merely as a technical configuration task, but as a central governance responsibility. This perspective aligns with broader AI governance scholarship, which emphasizes transparency, fairness, and institutional oversight in high-stakes educational technologies [9].

### A. Governance as a Primary Design Requirement

Traditional educational AI pipelines typically configure privacy budgets independently for each model or analytics component, resulting in inconsistent protection levels and limited institutional accountability [2], [3]. The proposed framework reframes differential privacy as an institution-level governance resource rather than a model-specific parameter, aligning with governance principles that emphasize transparency, human oversight, and auditability in AI deployments [9], [6]. By encoding policies and thresholds within a dedicated governance layer, institutions can enforce coherent and consistent privacy practices across tutoring, assessment, and analytics services.

**B. Balancing Utility and Fairness Through Multi-Objective Optimization**

Existing differential privacy mechanisms typically prioritize technical accuracy while overlooking demographic disparities. Utility-first allocation strategies can inadvertently intensify inequities across student subgroups—an issue well-documented in the fairness literature [4], [10], [18].

The proposed framework embeds fairness constraints directly into the allocation process, ensuring that privacy decisions do not place disproportionate burdens on particular student groups. By doing so, it supports more equitable learning outcomes while preserving strong pedagogical utility, consistent with insights from multi-objective optimization research [10], [12].

**C. Explainability as an Institutional Requirement**

Educational institutions increasingly require transparent AI decisions, especially when systems affect grading, feedback, or student profiling. Black-box optimization methods cannot meet these requirements because they lack interpretable rationales or audit trails. The proposed explainability layer provides:

- rule-based rationales,
- counterfactual reasoning,
- governance activation logs, and
- structured audit outputs.

These capabilities fill gaps identified in prior DP scheduling and fairness-aware systems, which lacked governance-aligned transparency [13], [14], [17].

**D. Integration with Cloud-Based and Multi-Service Edu-AI Systems**

Because the framework relies on interpretable policy classes—such as rule-based structures and monotonic additive models [23]—it integrates seamlessly with modern cloud-based educational ecosystems while remaining computationally efficient

This makes it applicable for:

- nationwide learning analytics platforms,
- university LMS deployments,
- accreditation-driven assessment systems,
- multi-service adaptive learning environments.

**E. Advancing Research on DP Governance**

The proposed method addresses several gaps across DP, fairness, and governance research:

- No previous model jointly optimizes utility, fairness, and explainability.
- Fairness-oriented methods do not incorporate DP constraints [10], [18].
- DP scheduling studies are domain-agnostic and ignore pedagogical utility [17], [14].
- Governance frameworks outline principles but lack operational allocation mechanisms [9].

By integrating these dimensions, the framework advances the field toward transparent and policy-aligned DP governance in education.

**F. Practical Implications for Real Deployment**

Institutions deploying AI-enabled learning systems can use the framework to: (i) implement principled and auditable privacy-budget governance, (ii) promote equitable outcomes across demographic groups, (iii) meet oversight expectations from ministries and accreditation bodies, and (iv) unify differential privacy, fairness, and explainability within a single governance architecture. Together, these capabilities position the framework as a practical blueprint for building trustworthy educational AI infrastructures.

### 7.1 Governance Risks and Limitations

Although the framework offers clear governance advantages, it also has important limitations. Mis-specified institutional rules or inadequately defined pedagogical utility metrics could lead to biased or suboptimal privacy allocations, highlighting the need for ongoing human oversight and periodic policy review. Utility-driven signals may also be vulnerable to strategic manipulation if incentive structures are not carefully designed. In addition, tensions may emerge between fairness objectives and institutional or regulatory mandates, especially in large-scale or cross-institutional deployments. Finally, while the framework is intentionally lightweight, scaling it to national or multi-institutional ecosystems will require coordinated governance policies and robust audit mechanisms. These considerations point to key avenues for future refinement and empirical validation.

## 8. Conclusion

This paper introduced a governance-aligned framework for explainable and fairness-aware differential privacy (DP) allocation in multi-service educational AI systems. In contrast to existing DP deployments—where privacy budgets are assigned independently and with minimal institutional oversight [2], [3]—the proposed approach reframes the global privacy budget as a strategic governance asset. Through a multi-objective optimization process, the framework aligns pedagogical utility, demographic fairness, and explainable governance, reflecting priorities emphasized in contemporary AI governance research [9].

### 8.1 The framework makes three primary contributions.

First, it offers a privacy-budget allocation model that jointly optimizes educational utility and subgroup fairness, drawing on principles from fairness-aware optimization literature [10], [12], [18]. Second, it introduces an interpretable allocation policy—implemented through rule-based or monotonic additive models [23]—that produces transparent, auditable decisions suitable for ministries, universities, and accreditation bodies. Third, it provides an explainability layer that generates governance logs, rule-based rationales, and counterfactual insights, addressing transparency gaps identified in DP scheduling and governance studies [13], [14], [17]. A conceptual comparison with standard baselines—including uniform allocation, utility-only strategies, and black-box optimizers—shows that the proposed framework uniquely achieves high utility, improved fairness, and full explainability. Existing approaches fail to satisfy all three requirements simultaneously, underscoring the need for a governance-oriented DP allocation mechanism in educational AI. Overall, the framework shifts differential privacy from a localized technical safeguard to an institution-level governance instrument, supporting the development of trustworthy, equitable, and accountable educational AI ecosystems.

### 8.2 Future Work

Future research should explore:

- •Empirical validation using real datasets such as OULAD or OpenEdu,
- •Dynamic real-time redistribution of privacy budgets as learner behavior evolves,
- •Human-in-the-loop governance mechanisms for institutional oversight, and
- •Scaling the governance layer for national or multi-campus cloud deployments.

These directions will strengthen the framework's readiness for large-scale, policy-aligned Edu-AI systems.

## References


1. N. Papernot et al., "Tempered sigmoid activations for deep differential privacy learning," in Proc. NeurIPS, 2022.

2. T. Yu and J. Li, "Differentially private adaptation in personalized education systems," IEEE Trans. Learning Technologies, vol. 16, no. 3, pp. 625–639, 2023.

3. J. Zhang and L. Chen, "Privacy-preserving student modeling with differential privacy," Computers & Education, vol. 204, 104886, 2024.

4. M. Hardt, E. Price, and N. Srebro, "Equality of opportunity in supervised learning," in Proc. NeurIPS, 2016, pp. 3315–3323.
5. Y. Li, S. Garg, and R. Buyya, "Fair resource allocation for AI-driven cloud systems," ACM Computing Surveys, 2023.
6. M. Alghamdi and D. Alanzi, "Readiness of institutions for adopting artificial intelligence applications," Arab Journal of Scientific Publishing, 2021.
7. L. Rosenblatt, J. Allen, and J. Stoyanovich, "Spending privacy budget fairly and wisely," arXiv preprint arXiv:2201.04890, 2022.
8. J. Hong, Z. Wang, and J. Zhou, "Dynamic privacy budget allocation improves model utility," in Proc. ACM FAccT, 2021.
9. S. Barocas, M. Madaio, and H. Wallach, "Fairness and machine learning in practice," in Proc. ACM FAccT, 2023.
10. F. Kamiran, A. Karim, and X. Zhang, "Discrimination-aware multi-objective model optimization," IEEE Trans. Knowledge and Data Engineering, vol. 34, no. 7, pp. 3451–3465, 2022.
11. T. Alabi and K. Wiggins, "Fairness-oriented privacy budget allocation for multi-user workloads," in Proc. ACM SoCC, 2023.
12. G. Yu, L. Ma, W. Du, and Y. Jin, "Towards fairness-aware multi-objective optimization," Complex & Intelligent Systems, 2025.
13. Q. Liu, R. Shakya, M. Khalil, and J. Jovanović, "Dynamic privacy budget controllers for machine learning systems," Journal of Machine Learning Research, 2025.
14. H. Elzayn, K. Ligett, S. Neel, and A. Roth, "Fair algorithms for learning in allocation problems," in Proc. ACM EC, 2019.
15. A. Alghamdi and P. Kalnis, "Fairness in resource allocation," IEEE Trans. Knowledge and Data Engineering, 2022.
16. X. Luo, Y. Zhang, and L. Wang, "Adaptive privacy-budget scheduling for data analytics," in Proc. IEEE ICDE, 2025.
17. F. Tholoniat, J. Smith, and K. Lee, "Hierarchical budget allocation for privacy-preserving systems," ACM Trans. Privacy and Security, 2024.

18. J. Kremer et al., “Fairness in multi-task learning,” Artificial Intelligence Review, 2024.
19. H. M. H. El-Hoby and A. A. Ibrahim, “Effective utilization of ICT resources in Saudi higher institutes,” European Journal of Technology, vol. 2, no. 2, pp. 1–16, 2017.
20. H. Mohamed, A. Abubakar, and A. Zeki, “Critical factors influencing effective ICT capabilities,” Journal of Theoretical and Applied Information Technology, vol. 96, no. 6, pp. 1–10, 2018.
21. H. M. H. El-Hoby, M. A. Salah, and M. A. Suhaimi, “Aligning cloud computing security with business strategy,” Int. J. Computer Technology, vol. 7, no. 1, pp. 52–60, 2014.
22. Y. Bengio, A. Lodi, and A. Prouvost, “Machine learning for combinatorial optimization: A methodological tour d’horizon,” *European Journal of Operational Research*, vol. 290, no. 2, pp. 405–421, 2021.
23. S. Meisenbacher, C. J. Lee, and F. Matthes, “Spend your budget wisely: Budget allocation in differential privacy,” in Proc. CODASPY, 2025.