Journal of Contemporary Science and Technology Management
Volume 02, Issue 04, December 2026

A Privacy-Preserving Blockchain and Federated Deep Learning Framework for Secure Healthcare Data Analytics

Manuel Verl
Uppsala, Sweden
Dr. Sandeep Kumar
Tulas University Dehradun
Sandeepkumarsiet@gmail.com

Abstract

Healthcare analytics increasingly relies on information dispersed across hospitals, laboratories and clinical networks. Centralized data pooling may expose sensitive records, create governance conflicts and undermine institutional autonomy. This study proposes a permissioned-blockchain and federated deep learning architecture that coordinates collaborative clinical prediction without transferring patient-level records. Each hospital trains a neural classifier locally; clipped, differentially private updates are protected by secure aggregation, while a permissioned ledger records authorization, model provenance and verifiable commitments rather than patient information. A threat model addresses curious coordinators, malicious participants, linkage attacks and ledger metadata leakage. The methodology specifies access policies, optimization objectives, secure update exchange, governance and reproducible evaluation protocols. An illustrative numerical scenario with five institutional nodes compares local-only, centralized, conventional federated and proposed protected federated configurations. The constructed example indicates that privacy and auditability can be strengthened at the expense of modest predictive and runtime overhead. These figures are scenario assumptions, not measured clinical findings. The framework offers a testable foundation for privacy-aware interinstitutional learning, with clinical deployment contingent on independent validation, security assessment and regulatory review.

Keywords: Federated deep learning; permissioned blockchain; healthcare analytics; differential privacy; secure aggregation; access control; auditability.

DOI: 10.66592/vgh8w123

1. Introduction

Electronic health records, imaging systems, laboratory platforms and wearable sensors create rich opportunities for clinical prediction. However, these datasets are fragmented across independent organizations. Institutional boundaries reflect patient consent, contractual obligations, differing technical infrastructures and the legitimate need to protect sensitive information. A model developed using a single hospital may fail when deployed in another setting because case mix, coding practice, device configuration and demographic composition differ. Collaboration is therefore scientifically desirable, but indiscriminate pooling of records is neither a universal technical solution nor an acceptable governance assumption.

Federated learning provides an alternative paradigm in which a coordinating service distributes a model to participating institutions, each institution updates it using its own records, and an aggregation procedure combines local contributions. The original federated averaging approach demonstrated that decentralized optimization can reduce the need to move raw observations [1]. In healthcare, federated training can enlarge the effective training population while allowing local custodians to retain physical control of their datasets [2]. Nevertheless, retaining records locally does not automatically ensure confidentiality. Gradients and model updates can encode information about individual examples, especially in small cohorts or under adversarial querying.

Privacy must therefore be treated as a system property rather than as a synonym for data localization. Stronger protections may include clipping, differential privacy, secure multiparty aggregation, encryption in transit, access minimization and rigorous participant authentication. Each mechanism addresses a different risk. Differential privacy constrains the influence of individual records on observable outputs; secure aggregation conceals individual institutional updates from the coordinator; encryption protects messages on the network; and institutional governance specifies which processing purposes are permissible. None of these measures by itself guarantees complete security against all adversaries.

A second challenge is accountability. Conventional federated learning commonly depends on a central server to schedule rounds, distribute model versions and record participation. If this service is compromised, disputed updates may be accepted, unauthorized parties may be enrolled, or model history may be altered. A permissioned blockchain can provide a shared, append-only record of selected operational events. The ledger is not a suitable repository for medical records or raw gradients: replication would increase privacy exposure and make correction or deletion difficult. Its defensible role is to record commitments, consent or authorization references, version identifiers and signed decisions.

The proposed framework combines a consortium ledger with privacy-enhanced federated deep learning. Hospital nodes act as independent data controllers or processors under an agreed governance arrangement. A membership service verifies institutional identities, smart contracts enforce eligibility rules, and a secure aggregation service computes global model updates. Only minimal metadata and cryptographic digests are committed to the ledger. Encrypted model artifacts and policy documents remain in governed off-chain storage. This separation limits unnecessary replication and allows data retention rules to be enforced at the storage layer.

Healthcare environments are heterogeneous. Some hospitals have thousands of observations, while others contribute smaller cohorts or serve specialized populations. The resulting non-identically distributed data can destabilize global optimization and produce unfair performance across sites. The system must evaluate site-level discrimination, calibration and subgroup differences instead of reporting global accuracy alone. A clinically useful model must be robust to missingness, coding changes, prevalence shift and temporal drift. Monitoring and the option to suspend a model are essential components of the proposed governance architecture.

This study asks whether an integrated design can improve the auditability of collaborative model development while maintaining practical prediction quality under explicit privacy controls. The central contribution is architectural and methodological: a complete round protocol, threat model, mathematical objective, ledger schema and evaluation plan are described. Because no real multicenter experiment was supplied, the numerical results are transparent illustrative scenarios rather than empirical evidence. This distinction is important: a fabricated benchmark cannot support claims of superior clinical effectiveness.

The manuscript makes four contributions. First, it identifies the privacy and governance boundaries between clinical records, model updates and ledger events. Second, it formulates weighted federated optimization with bounded updates and optional differential privacy. Third, it integrates authenticated participation, signed commitments and off-chain model custody into a coherent lifecycle. Fourth, it proposes a multidimensional assessment that includes prediction quality, privacy parameters, network cost, consensus latency and resilience to unauthorized requests. Together these elements establish a research design that can be reproduced and falsified in future controlled studies.

Clinical adoption also requires legal and ethical review. Local data protection obligations, institutional review board approval where applicable, patient information practices and purpose limitation remain necessary regardless of technical safeguards. A blockchain receipt is not equivalent to consent, and a privacy budget is not a substitute for evaluating reidentification risk. Security claims must be matched to realistic assumptions about collusion, compromised devices, identity management and operational procedures. The architecture therefore includes independent auditing and incident-response mechanisms rather than relying solely on cryptography.

Finally, the proposed approach should be compared with simpler alternatives. A centralized model may outperform federated alternatives when lawful data pooling is possible, while a conventional federated model may have lower

overhead where consortium governance is already trusted. The integrated framework is justified only if its additional traceability and policy enforcement deliver benefits that outweigh complexity and latency. The following sections review the literature, present a detailed methodology and provide a numerical illustration to support future experimental planning.

2. Literature Review

McMahan and colleagues introduced federated averaging as a communication-efficient strategy for optimizing deep networks using distributed client datasets [1]. Their formulation established the practical foundation for modern cross-device and cross-silo learning. For healthcare consortia, the central lesson is that raw-record centralization is not a prerequisite for model training. However, their optimization formulation does not itself provide a complete clinical governance, privacy accounting or immutable audit mechanism.

Rieke et al. surveyed federated learning in digital health and emphasized the promise of cross-institutional modeling alongside difficult issues of privacy, data heterogeneity, interoperability and reproducibility [2]. Their analysis is particularly relevant because hospitals differ from consumer devices in computing resources, organizational responsibilities and clinical safety requirements. They also cautioned that distributed training does not eliminate leakage risks. This work motivates the multilayer privacy design adopted here.

Kairouz et al. examined open problems in federated learning, including statistical heterogeneity, communication constraints, trust assumptions, personalization and secure computation [3]. These concerns imply that evaluation must move beyond a single aggregate accuracy score. For clinical settings, participation imbalance can distort training, and unequal site-level performance may cause downstream harms. Accordingly, the proposed evaluation reports both global and site-specific outcomes.

Dwork and Roth formalized differential privacy as a rigorous framework for bounding the effect of individual data contributions [4]. In learning systems, bounded gradients and calibrated noise can reduce information leakage, but privacy guarantees depend on the sampling scheme, clipping rule, accountant and repeated release of outputs. The present design treats the privacy budget as a measurable system parameter and avoids equating a low numerical epsilon with comprehensive institutional security.

Bonawitz et al. proposed practical secure aggregation to enable a server to recover a sum of client updates without inspecting each update individually [5]. This mechanism is valuable in federated healthcare because the coordinator may be operationally necessary but should not be able to inspect a particular hospital's contribution. Secure aggregation is not an antidote to malicious data poisoning, and collusion thresholds must be stated. The methodology therefore combines secure aggregation with admission controls and robust operational checks.

Androulaki et al. described Hyperledger Fabric as a modular permissioned blockchain with membership management and configurable transaction execution [6]. Such platforms are more suitable than public anonymous ledgers for institutional collaborations that require identifiable participants, access restrictions and predictable governance. Nevertheless, permissioned consensus introduces operational dependencies and metadata exposure. The ledger in this study stores minimal commitments rather than patient records or unprotected model updates.

Sheller et al. demonstrated collaborative learning across medical institutions without exchanging patient data in a brain-tumor segmentation context [7]. Their work illustrates how distributed medical imaging collaborations can be organized. However, the imaging setting differs from structured electronic health record prediction, and results cannot be transferred directly between tasks. The proposed framework is modality-agnostic in governance but specifies a tabular classification example for its evaluation design.

Li et al. investigated federated optimization in heterogeneous networks through FedProx [8]. Their contribution is relevant when sites differ in workload and statistical distribution. Proximal regularization can limit excessive divergence

from the global model, although its effectiveness depends on tuning and the nature of the shift. The present study includes an optional proximal term to accommodate hospital-specific case mixes and uneven local training.

Existing blockchain-for-healthcare proposals often emphasize record sharing, integrity or access management. Yet recording personal information directly on a replicated ledger conflicts with data minimization and correction requirements. A defensible design distinguishes evidence of an authorized operation from the sensitive operation itself. Signed hashes can help verify that an off-chain artifact matches an earlier commitment, but a hash alone does not establish that the artifact was clinically valid, lawfully processed or free from bias.

The literature therefore reveals a systems-level gap: federated optimization, privacy protection and distributed accountability are frequently evaluated separately. A complete design must specify exactly what is shared, what remains local, which adversary is resisted, how privacy is accounted for, and which metrics demonstrate operational value. The proposed contribution is not a claim that blockchain inherently improves prediction accuracy. Rather, it is a governance-aware integration that can be compared against simpler alternatives through controlled ablation studies.

3. Research Methodology

3.1 Research design and participating institutions

The study specifies a cross-silo consortium of five illustrative healthcare institutions, each holding its own deidentified or pseudonymized records in a secured local environment. The target task is binary prediction of a prespecified clinical outcome from structured variables such as age band, laboratory measurements and prior diagnosis indicators. An approved data dictionary standardizes units, missing-value indicators, time windows and outcome definitions. No patient-level rows are uploaded. The consortium maintains a shared protocol, institutional approval records and a pre-registered evaluation plan.

Each site splits its records by patient into training, validation and held-out test partitions before collaborative optimization. Temporal or external validation is preferred when available. Global metrics are calculated from privacy-preserving aggregated counts or locally computed summary statistics. All reported outcomes should include confidence intervals, class prevalence and calibration; site-level statistics are reviewed before any deployment decision.

3.2 Threat model and trust boundaries

The threat model includes an honest-but-curious aggregation coordinator, an unauthorized outsider, compromised hospital credentials, inference from released gradients, and a minority of malicious consortium participants. Secure aggregation protects individual updates from a noncolluding coordinator under its specified dropout and collusion assumptions. Differential privacy constrains record-level leakage under a stated accountant. Permissioned consensus deters unilateral rewriting of accepted event history but does not guarantee honest inputs or prevent endpoint compromise.

Journal of Contemporary Science and Technology Management
Volume 02, Issue 04, December 2026

3.3 System workflow

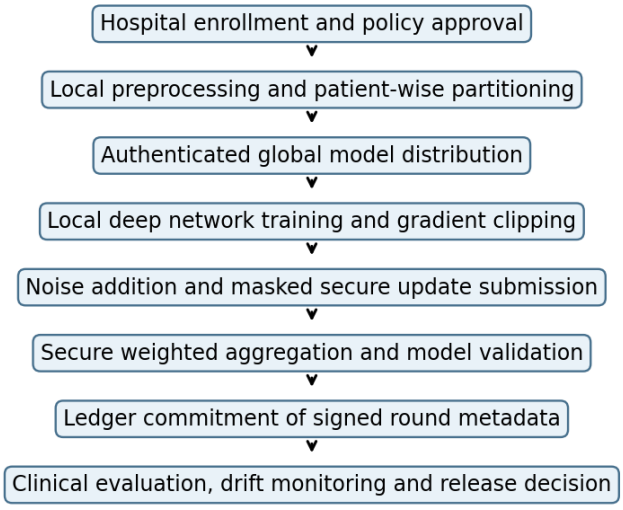

Figure 1. Proposed secure federated learning lifecycle (author-designed conceptual flowchart).

3.4 Local model and learning objective

Let K denote the number of institutions, n(k) the number of training records at institution k, and N the total training records. A multilayer perceptron with two hidden layers, rectified linear activations and sigmoid output is used for the illustrative binary prediction task. The weighted consortium objective is:

$$F(w) = \sum(k = 1\ to\ K)\ [n(k)/N]\ F(k, w)$$

Here w is the shared model parameter vector and F(k,w) is the empirical local loss. Each site uses weighted binary cross entropy to accommodate class imbalance:

$$F(k, w) = -(1/n(k))\sum(i = 1\ to\ n(k))[a\ y(i)\ ln\ p(i) + b(1 - y(i))ln(1 - p(i))]$$

The positive and negative class weights a and b are determined using training partitions only; y(i) is the observed binary outcome and p(i) is its predicted probability. Model hyperparameters are chosen using site-held validation sets and documented before the held-out evaluation.

3.5 Local optimization and heterogeneity

During communication round t, each eligible institution receives the signed global model version, trains for E local epochs and produces a bounded parameter update. A proximal penalty can limit local drift when site distributions are heterogeneous:

$$L(k, w) = F(k, w) + (\mu/2)\|w - w(t)\|^2$$

The nonnegative coefficient μ determines the strength of the proximal constraint; w(t) is the global parameter vector at the beginning of the round. When μ equals zero, the objective reduces to ordinary local empirical training. The baseline aggregate is a weighted mean:

$$w(t + 1) = \sum(k \in S(t))\ [n(k)/\sum(j \in S(t))n(j)]\ w(k, t + 1)$$

S(t) is the set of eligible sites completing the round. The weights are agreed before secure aggregation so that the coordinator cannot arbitrarily favor a participant. Dropout handling and minimum-participant thresholds are enforced before releasing any aggregate.

3.6 Differential privacy and secure aggregation

Each local update vector is clipped to a public Euclidean norm bound C, reducing the maximum influence of unusually large updates. The site then adds Gaussian noise calibrated to the selected privacy mechanism and accountant:

$$u(k) = \Delta w(k) \cdot min(1, C/\|\Delta w(k)\|)$$

$$z(k) = u(k) + \xi(k), \;\; \xi(k) \sim N(0, \sigma^2 C^2 I)$$

Here Δw(k) is the raw model update, u(k) is the clipped update, ξ(k) is a Gaussian noise vector, σ is the noise multiplier, and I is the identity matrix. These expressions define the perturbation operation, not a complete end-to-end differential privacy guarantee. The achieved privacy budget (ε,δ) must be calculated with a valid accountant using the actual sampling, participation, composition and release schedule. Institutional privacy may require additional protection beyond record-level privacy.

The secure aggregation protocol masks each weighted z(k) before transmission. The coordinator learns the weighted aggregate only if the required threshold of institutions completes the protocol. Transport-layer encryption, ephemeral keys, authenticated channels and replay protection complement masking. A model release gate blocks publication if the privacy budget is exhausted or the minimum participation threshold is not met.

3.7 Permissioned ledger and smart-contract governance

The consortium uses a membership-controlled ledger with roles for hospital trainer, aggregator, auditor and clinical approver. Smart contracts check participant status, protocol version, signed update receipt, privacy budget authorization and round state. The ledger records institution pseudonyms, timestamps, model version identifiers, content hashes and approval outcomes. Patient identifiers, clinical features, model gradients and encryption keys are never written on-chain.

$$c(t) = H(v(t) \;\|\; r(t) \;\|\; h(t) \;\|\; \tau(t))$$

H is a cryptographic hash function, v(t) the model version, r(t) the round identifier, h(t) the hash of the encrypted off-chain artifact, and τ(t) the timestamp. The symbol ‖ denotes concatenation. This commitment detects subsequent modification of a referenced artifact but does not verify its medical correctness.

The state machine progresses through Proposed, Authorized, Collecting, Aggregated, Validated and Finalized states. Invalid signatures, expired approvals or out-of-budget releases transition to Rejected. The consortium establishes key rotation, certificate revocation, dispute handling, audit retention and emergency shutdown procedures.

3.8 Reproducible evaluation protocol

A future empirical implementation should use a publicly documented clinical dataset with institutional partitions generated by site or a defensible non-IID simulation. Five random seeds, identical held-out cohorts and a fixed tuning budget are required for fair comparison. Baselines are local-only training, centralized pooling where legally and technically feasible, standard FedAvg, and FedAvg with privacy protections but without blockchain. The proposed architecture adds governance and ledger logging to the privacy-protected model.

The primary metric is area under the receiver operating characteristic curve (AUROC). Secondary metrics are accuracy, macro F1, sensitivity, specificity, expected calibration error, round duration, communication payload, ledger transaction latency and rejected unauthorized requests. Statistical uncertainty should be reported using patient-level bootstrap

Journal of Contemporary Science and Technology Management
Volume 02, Issue 04, December 2026

confidence intervals and site-stratified analyses. AUC and F1 are not interchangeable, and thresholds must be selected on validation data.

$$Accuracy = (TP + TN)/(TP + TN + FP + FN)$$

$$Precision = TP/(TP + FP),\ \ Recall = TP/(TP + FN)$$

$$F1 = 2 \cdot Precision \cdot Recall/(Precision + Recall)$$

TP, TN, FP and FN are the true-positive, true-negative, false-positive and false-negative counts. Privacy cost is recorded as the pair (ε,δ), not as a percentage of records protected. Security testing includes invalid-certificate requests, replayed transactions, unauthorized role actions, collusion scenarios and simulated poisoning. Attack success and defense coverage must be assessed separately.

3.9 Ablation, operational controls and limitations

The ablation plan removes ledger logging, secure aggregation, differential privacy and proximal regularization one at a time. This identifies which component contributes to privacy, governance, utility or computational overhead. The ledger should not be credited with improved predictive discrimination unless a mechanism and controlled result demonstrate such an effect. Robust aggregation and anomaly screening may improve resistance to poisoned updates, but they can conflict with secure aggregation because individual updates are hidden. The consortium must choose a compatible defense or evaluate trusted execution and verifiable computation alternatives.

Operational safeguards include secure hardware where feasible, independent penetration testing, artifact signing, role separation, data-quality checks, audit sampling and clinician review. Model cards document intended use, exclusions, prevalence, calibration and failure modes. Any deployment requires prospectively monitored performance and human oversight. Blockchain immutability complicates correction of erroneous metadata; corrections should therefore be appended as explicit superseding events, with sensitive records maintained in removable off-chain stores.

3.10 Illustrative configuration

Parameter	Illustrative setting
Institutions	5
Training records	50,000 total; uneven across sites
Prediction task	Binary clinical risk classification
Local epochs	3 per round
Communication rounds	40
Batch size	64
Clipping bound C	1.0
Noise multiplier σ	1.1
Privacy target	Accountant required; no verified ε claimed
Ledger	Permissioned consortium, off-chain model storage

These are proposed experimental settings. They do not indicate that any clinical dataset was processed, any blockchain was deployed or any privacy accountant was executed for this manuscript.

Journal of Contemporary Science and Technology Management
Volume 02, Issue 04, December 2026

4. Illustrative Results and Discussion

The following numerical tables are constructed scenario values for explaining the proposed experimental comparison. They are not observed results, cannot establish statistical significance and must be replaced with measurements before submission as an empirical research article.

Approach	AUROC	Accuracy	F1	Calibration error	Round time (s)
Local-only	0.824	0.781	0.756	0.064	18.0
Centralized	0.903	0.846	0.832	0.032	24.0
FedAvg	0.887	0.831	0.812	0.039	31.0
Private FL	0.872	0.818	0.799	0.043	37.0
Proposed FL + ledger	0.872	0.818	0.799	0.043	40.5

Table 1. Hypothetical comparative prediction and runtime outcomes. The ledger does not change the trained model in this scenario, so its predictive metrics match those of private federated learning. The added 3.5 seconds per round represents assumed governance and commitment overhead.

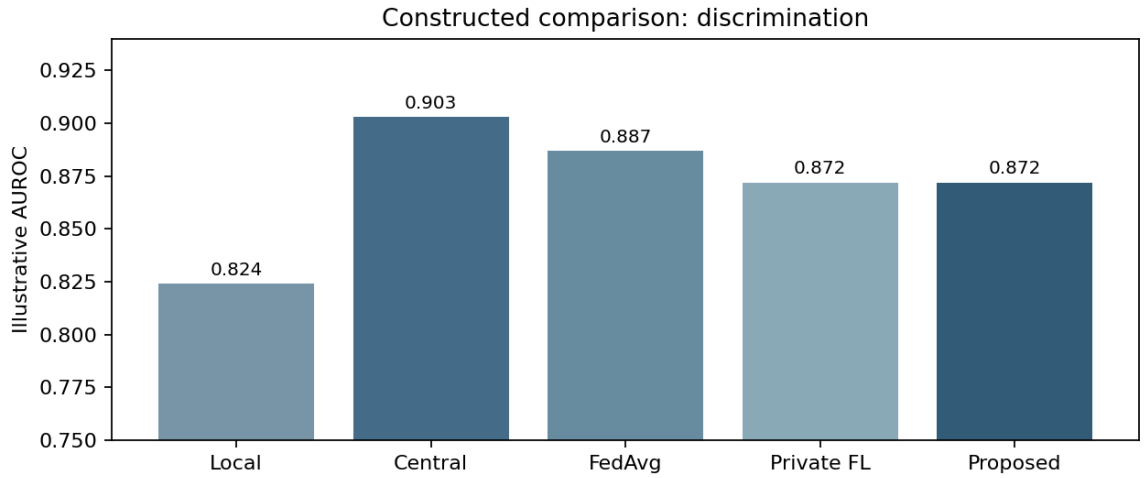

Figure 2. Illustrative AUROC values, not experimental measurements. In this scenario the private configuration sacrifices 0.015 AUROC relative to standard FedAvg, reflecting a possible privacy–utility trade-off rather than an established effect size.

Configuration	Round payload (MB)	Ledger latency (ms)	Unauthorized actions rejected (%)
FedAvg	42	Not applicable	Not assessed
Private FL	48	Not applicable	Not assessed
Proposed FL + ledger	48	210	99.0

Table 2. Hypothetical operational metrics. The 99% rejection figure is an illustrative assumption and is not evidence of a tested access-control system. Security outcomes should be reported with request counts, attacker capabilities and uncertainty.

Journal of Contemporary Science and Technology Management
Volume 02, Issue 04, December 2026

Noise multiplier	Illustrative AUROC	Illustrative round time (s)
0.0	0.887	35.0
0.6	0.880	39.0
1.1	0.872	40.5
1.5	0.858	41.5
2.0	0.839	43.0

Table 3. Constructed sensitivity scenario. Larger noise multipliers are shown with lower discrimination and slightly higher runtime. These values do not determine privacy budgets without a formal accountant and the full sampling protocol.

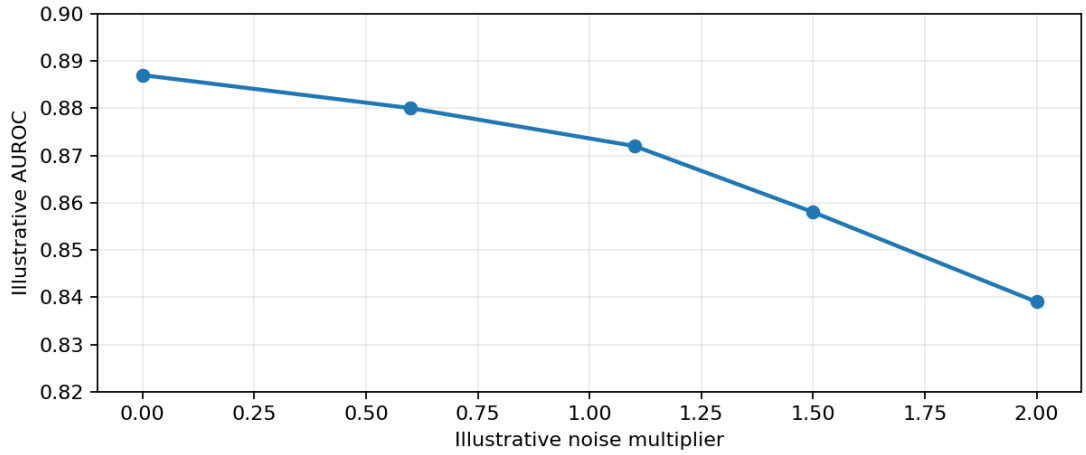

Figure 3. Constructed sensitivity curve illustrating a possible utility cost as noise increases. Real experiments may show different patterns depending on cohort size, task difficulty, optimization and accounting settings.

An additional governance ablation would compare a signed centralized log with the consortium ledger. Such a comparison is important because tamper-evident centralized logs may offer sufficient auditability with less complexity. The added value of distributed consensus should be measured through fault injection, dispute resolution, independent verification and operational cost rather than presumed from architecture alone.

5. Conclusion

This manuscript proposes an integrated privacy-preserving framework for collaborative healthcare analytics using federated deep learning, secure aggregation, differential privacy and a permissioned blockchain. Raw clinical records remain under institutional custody, while signed model commitments and policy decisions are recorded on a consortium ledger. The methodology specifies the optimization objective, update protection, governance lifecycle, threat model, evaluation metrics and deployment safeguards. The numerical illustrations demonstrate how utility, latency and accountability could be compared, but they are not empirical validation. Future work must implement the protocol, compute defensible privacy budgets, test adversarial behavior, measure site-level fairness and calibration, and conduct

independent clinical evaluation. The proposed design should be adopted only where its governance benefits justify the added system complexity.

References

Kaur, M., AlZubi, A. A., Jain, A., Singh, D., Yadav, V., & Alkhayyat, A. (2023). DSCNet: deep skip connections-based dense network for ALL diagnosis using peripheral blood smear images. Diagnostics, 13(17), 2752.

Devulapalli, N. (2019). Post-quantum cryptography for database security: Preparing for the impact of quantum computing on encrypted data storage and secure queries. Journal of Critical Reviews, 6(1), 614–620. https://doi.org/10.69980/jcr.v6:i1.13396

Jain, Jaishree, Arpit Jain, Saurabh Kumar Srivastava, Chaman Verma, Maria Simona Raboaca, and Zoltán Illés. "Improved security of e-healthcare images using hybridized robust zero-watermarking and hyper-chaotic system along with RSA." Mathematics 10, no. 7 (2022): 1071.

Kumar, S., Jain, A., & Moparthi, N. R. (2022). Enhanced method of object tracing using extended Kalman filter via binary search algorithm.

Khan, S. (2022). The use of artificial intelligence in ESG (environmental, social, and governance) investing: A study on AI models for sustainability analytics in asset and wealth management. International Journal of Multidisciplinary and Scientific Emerging Research, 10(1), 431–439. https://doi.org/10.15662/IJMSERH.2022.1001046

Talasila, D. (2023). Implementation of cloud-based disaster recovery models for minimizing business downtime and ensuring operational continuity through geographically distributed backup systems. International Journal on Recent and Innovation Trends in Computing and Communication, 11(8), 867–874. https://www.ijritcc.org/index.php/ijritcc/article/view/11934

Devulapalli, N. (2020). Implementation of compliance-as-code in regulated DevSecOps environments for automating audit readiness and reducing human error through scripted control validation. Journal of Critical Reviews, 7(1), 2651–2657. https://doi.org/10.69980/jcr.v7:i1.13405

Bhanushali, B. (2023). Blockchain for trade finance: Enhancing cross-border payment transparency, reducing processing time, and mitigating fraud through distributed ledger systems. International Journal on Recent and Innovation Trends in Computing and Communication, 11(9), 5690–5697. https://www.ijritcc.org/index.php/ijritcc/article/view/11932

Mahajan, R. A., Dey, R., Khan, M., Mohd Su'ud, M., Alam, M. M., & Jadhav, P. (2025). Enhancing personalization in IoT-based health monitoring via generative AI and transfer learning. Egyptian Informatics Journal, 32, Article 100788. https://doi.org/10.1016/j.eij.2025.100788

Devulapalli, N. (2020). The evolution of database security from centralized architectures to distributed cloud-native systems: A comparative and historical perspective. Turkish Online Journal of Qualitative Inquiry, 11(4), 2954–2963. https://doi.org/10.69980/j1064z97

Bhanushali, B. (2022). Blockchain-driven carbon credit trading systems for improving environmental transparency, reducing double counting, and ensuring accountability in climate finance. International Journal on Recent and Innovation Trends in Computing and Communication, 10(11), 292–300. https://www.ijritcc.org/index.php/ijritcc/article/view/11927

Rangappa, A. S. (2020). Deployment of cross-platform code scanning tools in polyglot DevSecOps projects for uniform security testing and language-agnostic vulnerability coverage. Journal of Critical Reviews, 7(11), 5000–5006. https://doi.org/10.69980/jcr.v7:i11.13402

Bhanushali, B. (2024). Tokenization of real-world assets through block chain for improving liquidity, enabling fractional ownership, and reducing transaction costs in global financial markets. International Journal of Advanced Research in Education and Technology, 11(4), 1809–1816. https://doi.org/10.15680/IJARETY.2024.1104071

Gupta, A. (2021). Deployment of advanced endpoint detection in remote work environments for preventing malware intrusions and data exfiltration through behavioral analytics and AI-driven responses. NeuroQuantology, 19(9), 1206–1224. https://doi.org/10.48047/nq.2021.19.9.NQ21199

Chatrath, A. (2021). Implementation of AI in recruitment platforms for improving candidate screening and diversity inclusion through bias-free resume parsing and skill matching. International Journal of Innovative Research in Computer and Communication Engineering, 9(12). https://doi.org/10.15680/IJIRCCE.2021.0912046

Anchala, S. (2023). Integration of confidential computing in multi-cloud environments for enhancing data privacy and regulatory compliance through encrypted processing and isolated execution. International Journal on Recent and Innovation Trends in Computing and Communication, 11(9), 5698–5709. https://www.ijritcc.org/index.php/ijritcc/article/view/11933

Chatrath, A. (2021). Integration of AI-powered forensic analysis in cloud breach investigations for accelerating root cause identification and evidence collection through log correlation techniques. International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering, 10(8). https://doi.org/10.15662/IJAREEIE.2021.1008051

Rangappa, A. S. (2020). Integration of behavioral analytics in developer workflows for detecting insider threats and malicious activities through activity monitoring and contextual profiling. International Journal of Innovative Research in Science, Engineering and Technology, 9(1), 13375–13383. https://doi.org/10.15680/IJIRSET.2020.0901064

Jain, Arpit, Adesh Kumar, and Sanjeev Sharma. "Comparative design and analysis of mesh, torus and ring NoC." Procedia Computer Science 48 (2015): 330–337.

Rangappa, A. S. (2021). Database security in cloud environments: Exploring multi-tenant risks, data isolation challenges, and compliance mechanisms in hybrid and public clouds. International Journal on Recent and Innovation Trends in Computing and Communication, 9(2), 50–57. https://www.ijritcc.org/index.php/ijritcc/article/view/11921

Jain, A. (2024, December). Energy-Efficient Signal Processing Architectures for IoT Networks Design. In 2024 International Conference on Emerging Technologies and Innovation for Sustainability (EmergIN) (pp. 362–367). IEEE.

Talasila, D. (2023). Enhancing cloud security with intrusion detection and prevention systems (IDPS): Comparative evaluation of signature-based, anomaly-based, and AI-powered detection models. International Journal on Recent and Innovation Trends in Computing and Communication, 11(3), 828–836. https://www.ijritcc.org/index.php/ijritcc/article/view/11945

Kumar, Ashok, and Arpit Jain. "Image smog restoration using oblique gradient profile prior and energy minimization." Frontiers of Computer Science 15, no. 6 (2021): 156706.

Chatrath, A. (2024). The role of encryption and key management in cloud security: Balancing performance, scalability, and confidentiality in multi-tenant and distributed data storage systems. International Journal of Innovative Research in Science, Engineering and Technology, 13(7). https://doi.org/10.15680/IJIRSET.2024.1307187

Jain, J., & Jain, A. (2022). Securing e-Healthcare Images Using an Efficient Image Encryption Model. Scientific Programming, 2022(1), 6438331.

Dwivedi, D. (2022). Application of AI-powered vulnerability scanners in legacy applications for identifying unpatched security flaws and weak dependencies through semantic code analysis. International Journal on Recent and Innovation Trends in Computing and Communication, 10(12), 617–623. https://www.ijritcc.org/index.php/ijritcc/article/view/11918

Dwivedi, D. (2020). Adoption of secure software updates in application lifecycle for preventing supply chain tampering and maintaining trust through cryptographically signed patches. Journal of Critical Reviews, 7(10), 6974–6980. https://doi.org/10.69980/jcr.v7:i10.13406

Jain, Arpit, Nageswara Rao Moparthi, A. Swathi, Yogesh Kumar Sharma, Nitin Mittal, Ahmed Alhussen, Zamil S. Alzamil, and MohdAnul Haq. "Deep Learning-Based Mask Identification System Using ResNet Transfer Learning Architecture." Computer Systems Science & Engineering 48, no. 2 (2024): 341.

Anchala, S. (2020). Cloud security challenges in supply chain systems: Addressing vendor risks, data integrity, and third-party service provider vulnerabilities. Journal of Critical Reviews, 7(2). https://doi.org/10.69980/jcr.v7:i2.13401

Talasila, D. (2020). Deployment of dynamic data masking in cloud-based applications for protecting sensitive information and enhancing privacy controls through contextual access filtering. Journal of Critical Reviews, 7(9), 4056–4062. https://doi.org/10.69980/jcr.v7:i9.13403

Kumar, S., Sharma, K., Kumar, P. A., Jain, A., Bhagat, S. K., & Singh, P. (2024, September). An improved particle swarm approach for energy-aware location-aided routing in mobile ad-hoc network. In 2024 7th International Conference on Contemporary Computing and Informatics (IC3I) (Vol. 7, pp. 1119–1124). IEEE.

Pandit, T., Mahendru, S., Raval, M., & Upadhyay, D. (2026). The evolution of reranking models in information retrieval: From heuristic methods to large language models. In X. Ding & Y. Dong (Eds.), Computational linguistics and natural language processing: CLNLP 2025 (Communications in Computer and Information Science, Vol. 2775, pp. 56–67). Springer. https://doi.org/10.1007/978-981-95-4788-3_6.

Gupta, A. (2021). Implementation of cyber threat hunting in hybrid cloud environments for early breach detection and threat containment through behavior-based analysis tools. NeuroQuantology, 19(12), 1312–1329. https://doi.org/10.48047/nq.2021.19.12.NQ21342

Gupta, A. (2021). Application of cyber deception strategies in military networks for enhancing operational security and misleading attackers through adaptive honeypots and decoy systems. NeuroQuantology, 19(4), 317–333. https://doi.org/10.48047/nq.2021.19.4.NQ21068

Dwivedi, D. (2019). The use of homomorphic encryption in cloud security: Enabling secure data processing, confidential query execution, and privacy-preserving cloud applications. International Journal of Innovative Research in Science, Engineering and Technology, 8(4), 4690–4698. https://doi.org/10.15680/IJIRSET.2019.0804136

Harshitha, G., Kumar, S., Rani, S., & Jain, A. (2021, November). Cotton disease detection based on deep learning techniques. In 4th Smart Cities Symposium (SCS 2021) (Vol. 2021, pp. 496–501). IET.

Ghori, M. M. The Future of Sustainable Packaging Systems: Architecting Resilient Manufacturing, Circular Supply Chains, and Regulatory Excellence. Global Pen Press UK.

S. R. Lankala, M. R. Marri, A. Jain, G. G. Battu, U. Lakhina and S. Singla, "Optimal Financial Fraud Detection and Alerting Mechanism in Cloud Computing Using Deep Belief Network," 2025 International Conference on Emerging Trends in Networks and Computer Communications (ETNCC), Windhoek, Namibia, 2025, pp. 743-748, doi: 10.1109/ETNCC66224.2025.11299665.

M. A. Arifin Shawn, N. Begum Asha, D. S. Jatav and R. Nair, "Federated Cloud Finance Ecosystem for Decentralized Risk Analytics and Sovereign Data Integrity through Cross-Institutional Learning," 2025 IEEE 7th International Conference on Cybernetics, Cognition and Machine Learning Applications (ICCCMLA), Hamburg, Germany, 2025, pp. 1-6, doi: 10.1109/ICCCMLA66092.2025.11580915

SinghJatav, D., Amin, M. M., Kodela, S., Nayan, V., Wannous, M., & Khalifa, G. S. (2025, November). Hybrid Reinforcement and Deep Learning Model for Payment Delay Optimization in Supply Chain Finance. In 2025 10th International Conference on Information Technology Trends (ITT) (pp. 170-175). IEEE.

V. Viswanathan, M. H. Mirza, D. S. Jatav, N. Mukhi, T. Gupta and S. B. Goyal, "Deep Reinforcement Learning Model to Enhance Accuracy of Forecasting in Supply chain Optimization," 2025 International Conference on Intelligent & Innovative Practices in Engineering & Management (IIPEM), Singapore, Singapore, 2025, pp. 1-6, doi: 10.1109/IIPEM65914.2025.11548310.

Kalal, M., & Tanner, O. C. (2025). Autonomous exception management in SAP S/4HANA manufacturing through multi-agent generative AI and event-driven supply networks. International Journal of Core Engineering & Management, 8(4), 130–137.

Kalal, M. (2026, February). Predictive Production Planning in Advanced Manufacturing Using SAP PP and Analytics. In 2026 5th International Conference on Sentiment Analysis and Deep Learning (ICSADL) (pp. 1363-1370). IEEE.

B. R. Rallabandi, K. Cherladine, N. Gudla and S. S. Mahadasu, ""AI-Enabled Private 5G and Non-Terrestrial Network Convergence for Sustainable Agricultural Connectivity,"" 2025 IEEE Connecting the Unconnected Europe, Middle East, and Africa Regional Summit (CTUS-EMEA), Abuja, Nigeria, 2025, pp. 1-9, doi: 10.1109/CTUS-EMEA68989.2025.11592222.

B. R. Rallabandi, K. Cherladine, S. S. Mahadasu and N. Gudla, ""Intent-Driven MVNO Architecture for Cross-Border 5G/6G Networks in Africa: Toward the One Africa Digital Economy,"" 2025 IEEE Connecting the Unconnected Europe, Middle East, and Africa Regional Summit (CTUS-EMEA), Abuja, Nigeria, 2025, pp. 1-11, doi: 10.1109/CTUS-EMEA68989.2025.11592181.

B. R. Rallabandi, D. Upadhyay, S. S. Mahadasu, D. S. Asudani, A. Jain and Y. K. Pandey, ""Investigation of Higher Order Sectorization Combined with Massive MIMO for Enhanced Spatial Diversity in 6G mmWave Environments,"" 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0, Raigarh, India, 2026, pp. 1-6, doi: 10.1109/OTCON69332.2026.11629684.