MANTECH Publications

Journal of Converging Technologies in Intelligent
Volume 1, Issue 2, May-August 2025

A Hybrid Deep Learning and Blockchain Framework for Secure and Decentralized Intelligent Data Management: A Review, Layered Architecture, and Simulation-Based Evaluation

Anshul K. Sarkar[1], Tanushree D. Mahato[2]

ABSTRACT

Intelligent data management in healthcare, industrial Internet of Things (IoT), and smart-city systems increasingly depends on deep learning, yet client–server training concentrates sensitive data, creates single points of failure, and offers no verifiable audit trail. Blockchain provides tamper-evident, decentralised record keeping but cannot analyse data or economically store large payloads. This review-based study synthesises representative work on deep learning, federated learning, and permissioned blockchains, identifies the remaining integration gaps, and proposes a three-layer hybrid framework in which edge nodes train deep models locally, encrypted payloads reside in distributed off-chain storage, and a permissioned ledger records content hashes, access policies, update validations, and participant reputation. A trust-weighted aggregation rule anchored on the ledger was evaluated by simulation against federated averaging and coordinate-wise median aggregation under model-poisoning attacks (20 clients, 40 rounds, 5 seeds). With 30% malicious clients, final test accuracy was 93.9 ± 3.2% for the proposed rule, 89.2 ± 5.7% for the median rule, and 38.5 ± 25.4% for federated averaging. Analytical models show an 8,192-fold reduction in on-chain storage through hash-anchoring and a 64-fold reduction in per-round consensus messages for crash-fault-tolerant ordering at 64 peers. The results support hybrid designs for secure, auditable, decentralised data management.

KEYWORDS: *Deep learning, Blockchain, Federated learning, Decentralized data management, Smart contracts, Data security, Byzantine fault tolerance, Model poisoning, Permissioned ledger, IPFS*

Page 133-146 © MANTECH PUBLICATIONS 2025. All Rights Reserved

INTRODUCTION

Data have become the principal input to modern decision-making, and deep learning has become the principal tool for extracting value from them [2]. Hospitals use convolutional networks to read radiological images, and utilities and municipal agencies use deep models to forecast load and detect anomalies in sensor streams. In nearly all of these deployments, however, raw data are copied to a central server where the model is trained. This arrangement concentrates sensitive records in a single administrative domain, exposes them to breach and insider misuse, and makes the resulting analytics difficult to audit, because the provenance of the training data and the integrity of the trained model cannot be independently verified by the parties that contributed them.

Two technological developments address complementary halves of this problem. Federated learning (FL) trains a shared model by exchanging parameter updates rather than raw data, so that data remain with their owners [3], [4]. Blockchain, introduced as the substrate for a decentralised currency [1], provides an append-only, cryptographically linked ledger that is maintained by mutually distrusting peers and is therefore tamper-evident without a trusted intermediary [5]. Neither is sufficient alone. Federated learning still relies on an aggregation server that can fail, collude, or be compromised, and its participants may submit poisoned updates. Blockchain, in contrast, offers neither analytical capability nor economical storage for large payloads, and public ledgers suffer from limited throughput and high energy cost [1], [5].

The central problem addressed in this paper is therefore how to combine deep learning and blockchain so that intelligent analytics can be performed over distributed, sensitive data while preserving confidentiality, integrity, availability, and accountability, without introducing prohibitive storage and communication overhead. The contributions are threefold: (i) a structured review of the integration literature; (ii) a three-layer hybrid framework with an explicit, ledger-anchored validation and reputation mechanism; and (iii) a simulation-based and analytical evaluation of its robustness and scalability.

LITERATURE REVIEW

Research on the intersection of artificial intelligence and blockchain has grown rapidly. Salah et al. classified the field into blockchain used to secure and decentralise AI workflows and AI used to improve blockchain operation, and identified data privacy, scalability, and resource

Page 133-146 © MANTECH PUBLICATIONS 2025. All Rights Reserved

cost as the principal open challenges [6]. Zheng et al. reached a similar conclusion for blockchain generally, listing scalability, privacy leakage, and selfish mining among its unresolved limitations [5]. Table 1 summarises representative frameworks that are discussed in the following subsections.

Deep Learning And Federated Learning

Deep neural networks learn hierarchical representations from large datasets and have achieved state-of-the-art results in vision, speech, and sequence modelling [2]. Their appetite for data conflicts with privacy regulation and institutional data-sharing reluctance. McMahan et al. proposed federated averaging (FedAvg), in which clients perform several local epochs of stochastic gradient descent and a server averages the resulting models weighted by local sample count, reducing communication rounds by one to two orders of magnitude relative to naive distributed SGD [3]. A broad survey of open problems identified statistical heterogeneity, communication cost, privacy leakage through gradients, and vulnerability to Byzantine participants as the main obstacles to deployment [4], [21]. In medicine, multi-institutional federated training for brain-tumour segmentation achieved accuracy approaching that of models trained on pooled data without moving patient images between sites [17]. Surveys of FL for IoT further show that resource-constrained devices and unreliable links impose additional design constraints [16].

Privacy And Robustness Mechanisms

Differential privacy bounds the influence of any individual record on a released statistic and can be applied to clipped model updates by adding calibrated noise [18]. Secure aggregation lets a server compute the sum of client updates without seeing any single update [19]. Robust aggregation addresses malicious participants: Krum selects the update closest to its neighbours and tolerates f Byzantine workers provided that the number of workers exceeds 2f + 2 [20]. These mechanisms are usually studied in a single-server setting, which leaves the aggregator itself as a trusted component.

Blockchain For Data Management And Learning

Early applications used the ledger as an access-control layer. Zyskind et al. stored encrypted personal data off-chain and kept permissions on-chain [11], and MedRec used smart contracts to manage electronic health record access across providers [12]. Both avoid placing bulk data

Page 133-146 © MANTECH PUBLICATIONS 2025. All Rights Reserved

on the ledger, a principle that is retained in the present work. The content-addressed InterPlanetary File System (IPFS) provides a natural off-chain store because an object's identifier is derived from its hash and can therefore be anchored on-chain for integrity verification [10]. Consensus choice determines both the fault model and the cost: practical Byzantine fault tolerance (PBFT) tolerates f faulty peers among $n \geq 3f + 1$ at a cost of quadratic message complexity [7], whereas Raft tolerates only crash faults but needs a linear number of messages per decision [9]. Permissioned platforms such as Hyperledger Fabric separate transaction execution from ordering and permit pluggable consensus [8].

Several studies combine the two technologies directly. BlockFL replaced the central FL server with a blockchain network and analysed the end-to-end latency created by block generation [13]. Lu et al. integrated a permissioned blockchain with federated learning for privacy-preserving data sharing in industrial IoT [14]. DeepChain introduced a blockchain-based incentive mechanism and auditability for collaborative deep learning [15]. Together, these works demonstrate feasibility but vary widely in fault model, validation logic, and cost treatment.

Table 1. Representative Blockchain–Learning and Blockchain–Data Frameworks

Framework	Ledger / Consensus	Learning or Data Role	Key Limitation
BlockFL [13]	Blockchain replaces the FL server	On-device federated learning	Block-generation latency
Lu et al. [14]	Permissioned blockchain	FL for industrial IoT data sharing	Consensus overhead on constrained nodes
DeepChain [15]	Blockchain-based incentive and audit	Collaborative deep learning	Cryptographic and incentive-layer overhead
MedRec [12]	Ethereum smart contracts	EHR access management (no learning)	No analytics; public-chain throughput
Zyskind et al. [11]	Blockchain access control, off-chain store	Personal-data protection (no learning)	No intelligent analytics layer

Page 133-146 © MANTECH PUBLICATIONS 2025. All Rights Reserved

RESEARCH GAP

The reviewed literature leaves three gaps. First, most blockchain-based FL designs treat the ledger as a passive log, whereas the validation of each update, and the resulting reputation of its contributor, could itself be a ledger-executed, auditable process [13], [15]. Second, the trade-off between fault model and consensus cost is seldom quantified for realistic consortium sizes, although it determines whether a design is deployable [7], [9]. Third, robustness to model poisoning and ledger-scale economics are generally reported separately, so that practitioners cannot judge the security gain and the infrastructure cost of a design together. This study addresses all three within one framework and one evaluation.

OBJECTIVES

- To review and compare representative deep learning, federated learning, and blockchain approaches for decentralised data management.
- To design a layered hybrid framework that keeps raw data and large payloads off-chain while anchoring integrity, access policy, and update validation on a permissioned ledger.
- To define a trust-weighted aggregation rule whose validation outcomes and reputation scores are recorded on the ledger.
- To evaluate the rule by simulation against federated averaging and median aggregation under increasing proportions of malicious clients.
- To quantify storage and consensus overhead analytically for healthcare, industrial IoT, and smart-city deployment scales.

METHODOLOGY

Review Protocol And Research Design

A structured narrative review was conducted, guided by the transparency principles of the PRISMA statement [22]. Works on federated learning, blockchain-based data management, consensus protocols, and robust aggregation were selected for foundational influence and relevance rather than exhaustive coverage. The research is therefore review-based, supplemented by simulation and analytical cost models; no clinical or field data were used.

Layered Framework Architecture

The proposed framework, shown in Figure 1, has three layers. In the edge layer, each data owner (for example a hospital, an industrial gateway, or a municipal data hub) trains a local

Page 133-146 © MANTECH PUBLICATIONS 2025. All Rights Reserved

deep model, clips the update, adds differential-privacy noise [18], and signs the result. In the off-chain layer, encrypted payloads and model checkpoints are held in IPFS-style distributed storage [10], and a secure-aggregation and analytics service computes the global model [19]. In the blockchain layer, a permissioned network of validator peers executes smart contracts that enforce access policy, validate updates, and record, for every round, the hash of each update, the storage content identifier (CID), the validation outcome, and the contributor's reputation. The ledger therefore holds evidence about data and models rather than the data and models themselves [11], [12].

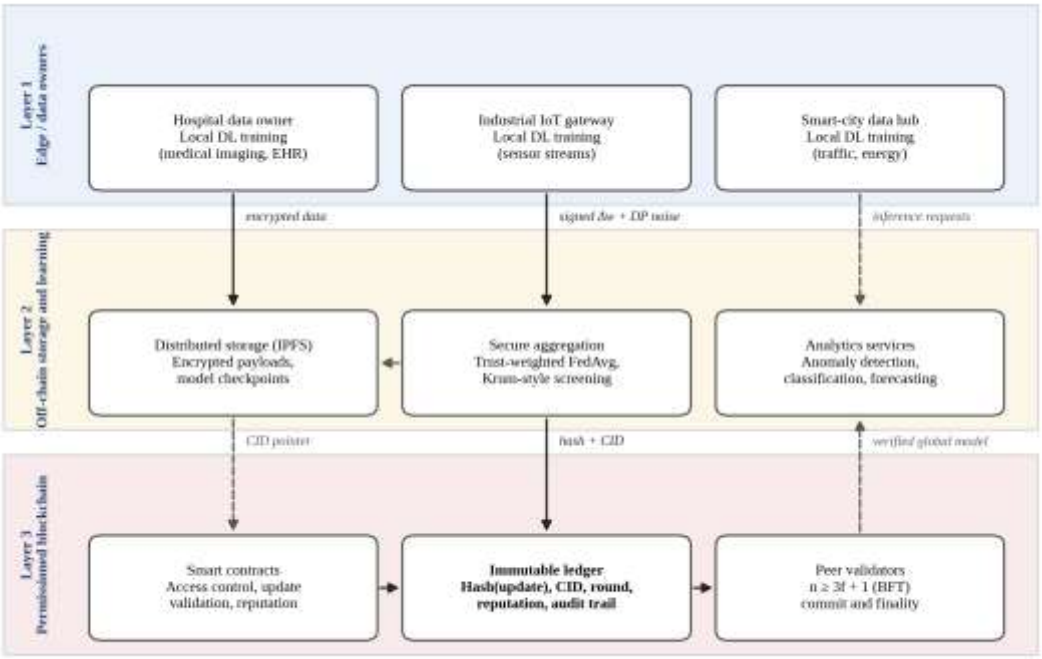


Figure 1. Proposed three-layer hybrid framework: edge learning, off-chain storage and secure aggregation, and a permissioned blockchain layer that anchors hashes, validation outcomes, and reputation.

Ledger-Anchored Trust-Weighted Aggregation

Let Δ_i denote the update of client i in a given round, n_i its local sample count, Δ_{med} the coordinate-wise median of all updates, and ρ_i its reputation. A smart contract computes a validity flag $a_i = 1$ when the cosine similarity between Δ_i and Δ_{med} is positive and $\|\Delta_i\| \leq 2 \times$

Page 133-146 © MANTECH PUBLICATIONS 2025. All Rights Reserved

median($\|\Delta_j\|$), and $a_i = 0$ otherwise. Reputation is updated as $\rho_i \leftarrow 0.8\rho_i + 0.2a_i$, and clients with $\rho_i < 0.3$ are excluded. The global model is then updated as:

$$w_{t+1} = w_t + \Sigma_i \left[n_i \rho_i a_i / \Sigma_j n_j \rho_j a_j \right] \Delta_i$$

Because the flags and reputations are written to the ledger by consensus, any participant can later verify why a contribution was accepted or rejected.

Simulation Set-Up

The simulation used a synthetic two-class dataset drawn from a Gaussian mixture in 20 dimensions (6,000 training and 4,000 test samples). Data were distributed over 20 clients with label skew generated from a Dirichlet distribution (concentration 0.8) to emulate non-identical local distributions [21]. Each client trained a multilayer perceptron (20–32–1, tanh hidden units) for two local epochs (learning rate 0.1, batch size 32) in each of 40 rounds. A proportion of clients (0, 10, 20, 30, or 40%) mounted a model-poisoning attack by submitting their update multiplied by −3. Three aggregators were compared: FedAvg, coordinate-wise median, and the proposed rule. Each configuration was repeated with five random seeds, and the final test accuracy is reported as mean ± standard deviation.

Analytical Cost Models

Per-round message counts were modelled as $2n(n-1)$ for PBFT and $2(n-1)$ for Raft, where n is the number of validator peers. On-chain storage was modelled assuming 2 MB per raw record versus a 256-byte ledger entry containing the content hash, CID, metadata, and signature. Model-update volume was computed for a 11.7-million-parameter network stored in 32-bit floating point (about 46.8 MB per update).

RESULTS AND FINDINGS

Figure 2 presents the analytical overhead results. Consensus cost diverges sharply with network size: at n = 64 peers, PBFT requires 8,064 messages per round whereas Raft requires 126, a 64-fold difference, and the gap widens linearly with n (Figure 2a). Hash-anchoring decouples ledger growth from payload size: for one million records of 2 MB each, a full on-chain design would require about 2 TB, whereas the anchored design requires about 244 MB, a reduction of 8,192-fold (Figure 2b). Table 2 translates these models to three deployment scales.

Page 133-146 © MANTECH PUBLICATIONS 2025. All Rights Reserved

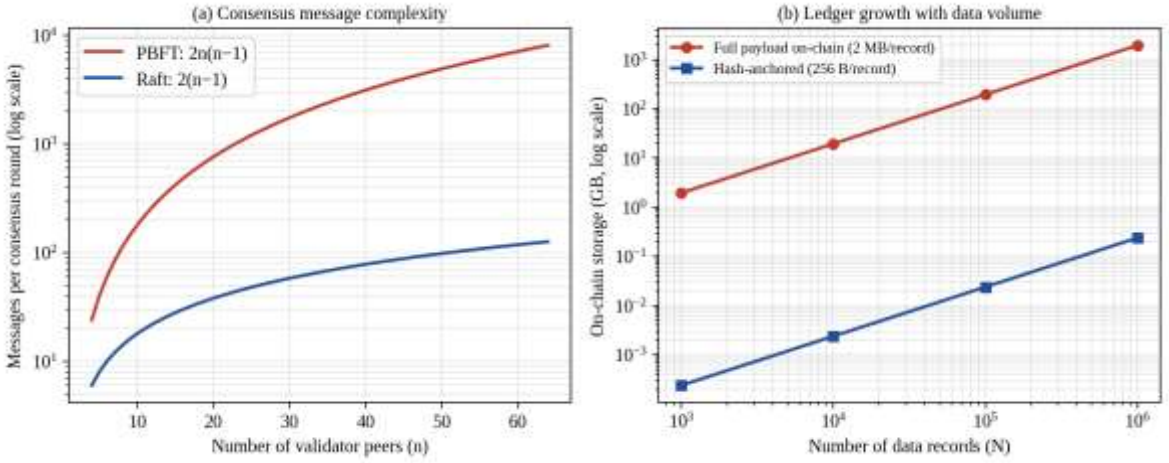


Figure 2. Analytical overhead: (a) consensus messages per round for PBFT and Raft versus number of peers; (b) on-chain storage versus number of records for full-payload and hash-anchored designs.

Table 2. Analytical Overhead of the Framework at Three Deployment Scales

Scenario	Peers (n)	PBFT Messages / Round	Raft Messages / Round	Ledger Growth / Round	Off-Chain Updates / Round
Healthcare consortium	8	112	14	2.0 KB	374 MB
Industrial IoT	32	1,984	62	8.0 KB	1.50 GB
Smart city	64	8,064	126	16.0 KB	2.99 GB

Figure 3 and Table 3 present the simulation results. Without attackers, the three aggregators performed comparably (94.7 ± 3.0% for FedAvg, 93.3 ± 3.3% for the median, 94.5 ± 3.3% for the proposed rule), so the validation step cost 0.2 percentage points relative to FedAvg. Under attack, FedAvg degraded rapidly: accuracy fell to 82.6 ± 11.0% at 20% malicious clients, 38.5 ± 25.4% at 30%, and 14.5 ± 7.3% at 40%, which is below chance because the scaled, sign-reversed updates actively drive the model away from the solution. The median rule resisted attack more effectively but lost accuracy progressively, reaching 89.2 ± 5.7% at 30% and 86.8 ± 4.8% at 40%. The proposed rule retained 93.9 ± 3.2% at 30% and 92.1 ± 5.2% at 40%, and at 30% malicious clients its mean accuracy exceeded 90% after only six rounds, whereas the median rule did not reach 90% in 40 rounds.

Page 133-146 © MANTECH PUBLICATIONS 2025. All Rights Reserved

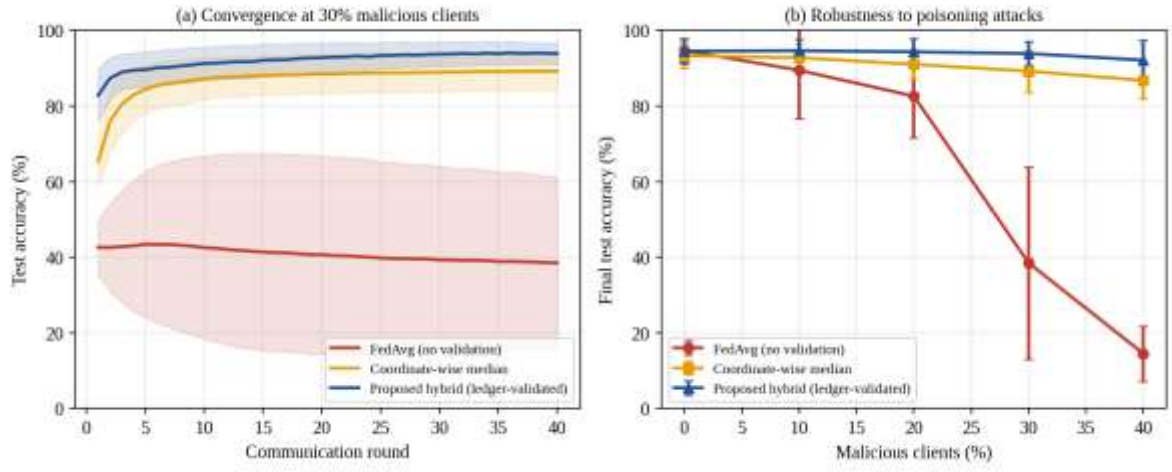


Figure 3. Simulation results: (a) test accuracy per round with 30% malicious clients; (b) final accuracy versus proportion of malicious clients (mean ± SD, 5 seeds).

Table 3. Final Test Accuracy (%) under Model-Poisoning Attack (mean ± SD, 5 seeds)

Malicious Clients	FedAvg	Coordinate-wise Median	Proposed Hybrid Rule
0%	94.7 ± 3.0	93.3 ± 3.3	94.5 ± 3.3
10%	89.4 ± 12.8	92.8 ± 3.7	94.6 ± 3.1
20%	82.6 ± 11.0	91.0 ± 3.9	94.3 ± 3.5
30%	38.5 ± 25.4	89.2 ± 5.7	93.9 ± 3.2
40%	14.5 ± 7.3	86.8 ± 4.8	92.1 ± 5.2

DISCUSSION

Interpretation Of Findings

The simulation confirms the intuition that averaging is a linear operator and is therefore hostage to its most extreme input [20]. The median rule bounds the influence of outliers but discards information from honest clients with legitimate, skewed data, which explains its gradual loss of accuracy. The proposed rule combines geometric screening within a round with memory across rounds: a persistent attacker accumulates low reputation and is excluded, whereas an honest client with an unusual round is penalised only transiently. This temporal dimension is only possible because reputation is durable state, and the ledger supplies that state without a trusted server.

Page 133-146 © MANTECH PUBLICATIONS 2025. All Rights Reserved

The analytical results show that the dominant cost drivers can be controlled by architecture rather than by faster hardware. Keeping payloads and model tensors off-chain reduces ledger volume by orders of magnitude, and choosing the consensus protocol to match the trust assumption determines whether messaging cost grows linearly or quadratically with consortium size [7], [9]. PBFT remains affordable for consortia of a few dozen mutually distrusting organisations; larger or more trusted federations benefit from crash-fault-tolerant ordering.

Engineering And Scientific Significance

From an engineering standpoint, the framework offers a reference decomposition in which each layer is replaceable: the learner (any federated deep model), the store (any content-addressed system), and the consensus engine (BFT or Raft). From a scientific standpoint, it frames reputation and validation as ledger-resident, auditable objects, which opens a path to formal accountability analysis of collaborative learning.

Practical Applications

In healthcare, hospitals could jointly train diagnostic models while the ledger records which institutions contributed to which model version, supporting consent and regulatory audit [12], [17]. In industrial IoT, equipment manufacturers and operators could share failure-prediction models without exposing proprietary process data [14]. In smart cities, transport, energy, and environmental agencies could build shared forecasting models, with the ledger providing a neutral record of contributions that is not owned by any single agency [16].

LIMITATIONS

- The simulation used synthetic tabular data and a small multilayer perceptron; behaviour on large convolutional or transformer models and on real, heterogeneous datasets has not been established.
- Only one attack family (scaled sign reversal) was tested; adaptive adversaries who craft updates to evade cosine and norm checks were not modelled.
- The ledger and consensus were modelled analytically rather than executed on a running permissioned network, so latency, throughput, and smart-contract execution cost were not measured.

Page 133-146 © MANTECH PUBLICATIONS 2025. All Rights Reserved

- The review is selective rather than exhaustive, and formal privacy guarantees depend on the noise level and privacy-budget composition over rounds [18].
- Reputation is vulnerable to Sybil identities unless the permissioned network enforces strong identity management.

FUTURE SCOPE

- Prototype implementation on a Hyperledger Fabric network with IPFS storage to measure end-to-end latency, throughput, and chaincode cost under realistic loads [8], [10].
- Evaluation against adaptive and backdoor attacks and against alternative robust aggregators, combined with secure aggregation so that individual updates are never exposed [19], [20].
- Formal analysis of the privacy budget across rounds, and exploration of zero-knowledge proofs that allow validation of an update without disclosing it.
- Incentive design in which reputation translates into rewards, and extension to cross-chain federations of consortia.

CONCLUSION

This study reviewed the integration of deep learning and blockchain, identified the absence of ledger-resident validation and combined cost–robustness evaluation, and proposed a three-layer framework that addresses both. Simulation showed that ledger-anchored, trust-weighted aggregation preserved test accuracy of 93.9% with 30% malicious clients, compared with 38.5% for federated averaging and 89.2% for median aggregation, at a validation cost of 0.2 percentage points in the benign case. Analytical models showed that off-chain storage with on-chain hashes reduces ledger growth by a factor of 8,192 for 2 MB records, and that consensus selection can reduce per-round messaging by more than 60-fold at 64 peers. Hybrid architectures of this kind are a credible basis for secure, auditable, and decentralised intelligent data management, subject to the validation on real networks and datasets outlined in the future scope.

REFERENCES

1. Nakamoto S. Bitcoin: A peer-to-peer electronic cash system. 2008. Available at: https://bitcoin.org/bitcoin.pdf
2. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature.* 2015; 521(7553): 436–444.

Page 133-146 © MANTECH PUBLICATIONS 2025. All Rights Reserved

3. McMahan B, Moore E, Ramage D, Hampson S, y Arcas BA. Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), PMLR.* 2017; 54: 1273–1282.
4. Kairouz P, McMahan HB, Avent B, et al. Advances and open problems in federated learning. *Foundations and Trends in Machine Learning.* 2021; 14(1–2): 1–210.
5. Zheng Z, Xie S, Dai H, Chen X, Wang H. Blockchain challenges and opportunities: a survey. *International Journal of Web and Grid Services.* 2018; 14(4): 352–375.
6. Salah K, Rehman MHU, Nizamuddin N, Al-Fuqaha A. Blockchain for AI: Review and open research challenges. *IEEE Access.* 2019; 7: 10127–10149.
7. Castro M, Liskov B. Practical Byzantine fault tolerance. *Proceedings of the 3rd Symposium on Operating Systems Design and Implementation (OSDI).* 1999; 173–186.
8. Androulaki E, Barger A, Bortnikov V, et al. Hyperledger Fabric: A distributed operating system for permissioned blockchains. *Proceedings of the 13th EuroSys Conference.* 2018; Article 30: 1–15.
9. Ongaro D, Ousterhout J. In search of an understandable consensus algorithm. *Proceedings of the USENIX Annual Technical Conference.* 2014; 305–319.
10. Benet J. IPFS – Content addressed, versioned, P2P file system. *arXiv preprint* arXiv:1407.3561. 2014.
11. Zyskind G, Nathan O, Pentland A. Decentralizing privacy: Using blockchain to protect personal data. *Proceedings of the IEEE Security and Privacy Workshops.* 2015; 180–184.
12. Azaria A, Ekblaw A, Vieira T, Lippman A. MedRec: Using blockchain for medical data access and permission management. *Proceedings of the 2nd International Conference on Open and Big Data (OBD).* 2016; 25–30.
13. Kim H, Park J, Bennis M, Kim SL. Blockchained on-device federated learning. *IEEE Communications Letters.* 2020; 24(6): 1279–1283.
14. Lu Y, Huang X, Dai Y, Maharjan S, Zhang Y. Blockchain and federated learning for privacy-preserved data sharing in industrial IoT. *IEEE Transactions on Industrial Informatics.* 2020; 16(6): 4177–4186.
15. Weng J, Weng J, Zhang J, Li M, Zhang Y, Luo W. DeepChain: Auditable and privacy-preserving deep learning with blockchain-based incentive. *IEEE Transactions on Dependable and Secure Computing.* 2021; 18(5): 2438–2455.

Page 133-146 © MANTECH PUBLICATIONS 2025. All Rights Reserved

16. Nguyen DC, Ding M, Pathirana PN, Seneviratne A, Li J, Poor HV. Federated learning for Internet of Things: A comprehensive survey. *IEEE Communications Surveys & Tutorials.* 2021; 23(3): 1622–1658.
17. Sheller MJ, Edwards B, Reina GA, et al. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Scientific Reports.* 2020; 10: 12598.
18. Dwork C, Roth A. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science.* 2014; 9(3–4): 211–407.
19. Bonawitz K, Ivanov V, Kreuter B, et al. Practical secure aggregation for privacy-preserving machine learning. *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS).* 2017; 1175–1191.
20. Blanchard P, El Mhamdi EM, Guerraoui R, Stainer J. Machine learning with adversaries: Byzantine tolerant gradient descent. *Advances in Neural Information Processing Systems.* 2017; 30: 119–129.
21. Li T, Sahu AK, Talwalkar A, Smith V. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine.* 2020; 37(3): 50–60.
22. Moher D, Liberati A, Tetzlaff J, Altman DG. Preferred reporting items for systematic reviews and meta-analyses: The PRISMA statement. *PLoS Medicine.* 2009; 6(7): e1000097.

Page 133-146 © MANTECH PUBLICATIONS 2025. All Rights Reserved

***Author for Correspondence:**
Anshul K. Sarkar
E-mail: anshul.k.sarkar456@outlook.com

[1]Assistant Professor, Department of Computer Science and Engineering, Heeralal Yadav Institute of Technology & Management, Lucknow
[2]Associate Professor, Department of Computer Science and Engineering, Heeralal Yadav Institute of Technology & Management, Lucknow

Received Date: May 23, 2025
Accepted Date: June 07, 2025
Published Date: June 22, 2025

Citation: Anshul K. Sarkar, Tanushree D. Mahato. A Hybrid Deep Learning and Blockchain Framework for Secure and Decentralized Intelligent Data Management: A Review, Layered Architecture, and Simulation-Based Evaluation. Journal of Converging Technologies in Intelligent. 2025; 1(2): 133-146p.

Page 133-146 © MANTECH PUBLICATIONS 2025. All Rights Reserved