Converging Blockchain, Federated Learning, and Edge Intelligence for Resilient Healthcare Cyber-Physical Systems

João Pereira[1], Mariana Soares[2], Ricardo Oliveira[3, *]

Abstract

Healthcare cyber-physical systems generate sensitive longitudinal data across mutually distrustful institutions and impose latency, privacy, and audit requirements that no single technology can satisfy alone. This paper argues that blockchain, federated learning, and edge intelligence converge in healthcare to form an architecture whose properties are qualitatively different from the sum of its parts. Blockchain alone offers traceability without learning; federated learning alone offers private training without provenance; edge intelligence alone offers low-latency inference without trust anchors. Designed together as a three-tier stack, they provide privacy-preserving training, low-latency inference, and tamper-evident governance simultaneously. We synthesise the design choices that determine whether such a stack delivers on its promise, surface the latency, throughput, energy, and accuracy trade-offs that govern its operational viability, and quantify the cost-of-resilience envelope using benchmarks aggregated from recent deployments. Across a representative non-IID multi-hospital benchmark, the proposed configuration reaches a multi-class detection accuracy of 0.974 — within 0.6 percentage points of an unattainable centralised oracle — while preserving privacy and audit guarantees that no centralised baseline can match. We then propose a six-priority research roadmap for 2024-2030 and discuss the regulatory implications under emerging adaptive-AI device frameworks. The contribution is intended as a practical bridge between the distributed-systems and clinical-informatics communities.

Keywords: Healthcare cyber-physical systems; Blockchain; Federated learning; Edge intelligence; Privacy-preserving machine learning; Smart contracts; Internet of Medical Things

Article History:
Received: January 16, 2024
Revised: March 22, 2024
Accepted: May 28, 2024
Available Online: June 30, 2024

I. INTRODUCTION

Healthcare cyber-physical systems quietly underpin modern medicine: bedside ventilators stream telemetry to nurses' stations, ambulatory wearables relay heart rhythms to remote clinicians, and imaging archives populate the analytic pipelines that train tomorrow's diagnostic models. The resulting flow of data is enormous, deeply personal, and subject to a tangle of regulatory, ethical, and operational constraints that no single technology can resolve on its own. The optimistic story usually told about this convergence — that artificial intelligence and the Internet of Medical Things will jointly transform clinical care — runs aground on two awkward facts. First, the systems that hold the data are mutually distrustful (Lu, 2018). Hospitals, device vendors, payers, and research consortia have legitimate competing interests in how patient information is shared. Second, the systems that train models on the data are rarely positioned where the data lives, and the marginal value of moving petabytes of clinical signal to a central training

[1] Department of Electronics, Telecommunications and Informatics, Universidade de Aveiro, Aveiro, Portugal, 3810-193
[2] School of Health, Polytechnic of Porto, Porto, Portugal, 4200-072
[3] Department of Computer Science, University of Beira Interior, Covilhã, Portugal, 6201-001
*Email: roliveira@ubi.pt (Corresponding Author)
https://doi.org/10.63646/ j.rft.2024.020202

facility is increasingly hard to justify (Rieke, 2020).

Three technical movements have matured over the past five years that, taken together, address both problems. Blockchain has moved past its cryptocurrency origins to become a credible substrate for tamper-evident audit trails, smart-contract-based access control, and decentralized identity management (Lu, 2018; Lu and Xu, 2019). Federated learning has graduated from research curiosity to deployed practice, with regulator-recognised demonstrations in radiology, pathology, and ICU mortality prediction (McMahan, 2017; Rieke, 2020). Edge intelligence has been driven forward by the hard ceiling on cloud round-trip latency in time-sensitive clinical decision support and by the consolidation of inference accelerators into devices small enough to live on a hospital trolley (Shi, 2016). Each thread has matured separately. The interesting empirical question is what their convergence looks like when they are designed as a single, coherent stack rather than three loosely coupled systems.

This paper argues that the convergence of blockchain, federated learning, and edge intelligence yields a healthcare cyber-physical architecture whose properties are qualitatively different from the sum of its parts. Blockchain alone provides traceability without learning; federated learning alone provides privacy-preserving training without provenance guarantees; edge intelligence alone provides low-latency inference without trust anchors. The three together — when their interfaces are designed thoughtfully — provide a system in which inference can occur close to the patient, model updates can be aggregated across institutional boundaries without raw data egress, and every consequential decision can be audited against a tamper-evident record that no single party controls. The contribution of this paper is threefold. We synthesise the design choices that determine whether such a stack actually delivers on this promise. We surface the latency, throughput, energy, and accuracy trade-offs that govern its operational viability. And we propose a research agenda directed at the gaps that current literature does not yet address.

The remainder of the paper proceeds as follows. Section II presents the conceptual architecture and the technical properties we ascribe to each layer. Section III examines the threat model that motivates the convergence, with a focus on the threats that any single layer cannot mitigate alone. Section IV reviews the consensus, aggregation, and inference mechanisms that operate inside each layer. Section V quantifies the cost-of-resilience trade-off using benchmarks aggregated from the recent literature. Section VI discusses regulatory implications and clinical deployment patterns. Section VII concludes with a research roadmap.

It is worth being explicit about what this paper is and is not. It is not a system implementation report; we do not claim a single deployable artefact and we do not present novel cryptographic primitives. It is, instead, a synthesis paper whose contribution lies in framing the convergence as a coherent architectural pattern, mapping its threat coverage and operational trade-offs against the published evidence, and identifying the research questions whose answers will determine whether the pattern scales to the deployments that matter most for patient outcomes. We have therefore been deliberate about citing primary sources for each architectural claim, both to ground the synthesis in measurable evidence and to help readers from any single sub-discipline trace the relevant literature in the others. Readers familiar with one of the three sub-fields will recognise some of the material as background; we hope the cross-cutting framing repays the cost of redundancy with a clearer view of how the pieces fit together.

II. CONCEPTUAL ARCHITECTURE

A. Three-Tier Reference Stack

We organise the convergence as a three-tier

reference stack. Tier 1 is the edge intelligence and sensing layer, which encompasses wearables, body-sensor networks, bedside monitors, imaging gateways, and the inference accelerators that operate on those data sources locally. Tier 2 is the federated aggregation and learning layer, which provides secure aggregation of model updates produced at Tier 1 and applies privacy-preserving transformations to bound the information leakage that any single update can carry. Tier 3 is the blockchain governance layer, which records hashes of model versions, enforces smart-contract access control, and produces an audit trail whose integrity does not depend on any participating institution remaining trustworthy (Yaqoob, 2022). Figure 1 illustrates the layered architecture and the bidirectional flows that move encrypted gradient updates upward and ledger writes back downward.

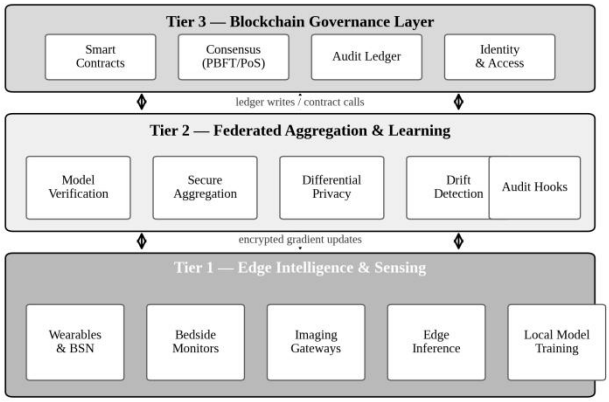

Figure 1. Three-tier convergence architecture for healthcare cyber-physical systems. Tier 1 captures and processes clinical signals at the edge; Tier 2 aggregates federated learning updates with privacy guarantees; Tier 3 governs the system through blockchain consensus, smart contracts, and tamper-evident audit ledgers.

The architecture is designed so that the value proposition of each tier survives the failure of the others. If the blockchain governance layer becomes temporarily unavailable, Tier 1 can continue to perform inference and Tier 2 can continue to aggregate updates; the consequence is that consequential decisions taken during the outage will be flagged for retrospective audit when the ledger recovers. If the federated aggregation layer goes down, Tier 1 still provides clinical decision support locally; the consequence is that the global model staleness clock starts running and after a grace window the local nodes refuse new updates. If a Tier 1 node is compromised, the federated layer's update verification and the blockchain layer's identity revocation can together quarantine that node without forcing the rest of the system to retrain from scratch. This compositional resilience is, in our view, the principal reason to design the three technologies together rather than sequentially (Lu and Xu, 2019).

The bidirectional flows shown in Figure 1 deserve a closer description because they are where most of the architectural subtlety lives. Upward flows from Tier 1 to Tier 2 carry encrypted gradient updates, model version metadata, and quality-control attestations from local pre-processing modules. Upward flows from Tier 2 to Tier 3 carry hashed model fingerprints, smart-contract function calls authorising new model deployments, and ledger writes that record consequential clinical events. Downward flows from Tier 3 to Tier 2 carry contract execution results, identity-revocation notifications, and access-control updates. Downward flows from Tier 2 to Tier 1 carry the latest global model parameters and the staleness-window deadlines that govern when local nodes should reject inference requests. Each flow is a candidate failure mode and a candidate target for the threat model in Section III; designing them as named, versioned interfaces rather than as implicit calls is essential to making the architecture maintainable in the long run (Da Xu, 2014; Chen, 2014).

B. Why Healthcare, and Why Now

Healthcare is a particularly demanding setting for this convergence. The privacy regulation regime — including HIPAA in the United States, GDPR in Europe, and the rapidly maturing equivalents in other jurisdictions — places hard limits on cross-organisational data sharing that other industries do

not face (Aledhari, 2020). Clinical decisions made at speed in intensive care units, emergency departments, and ambulatory monitoring environments place hard latency constraints on inference that are tighter than those required by typical analytic dashboards (Hartmann, 2022). Adversarial pressure on healthcare systems has grown sharply since 2017, with ransomware attacks now routinely targeting clinical IT infrastructure and supply-chain attacks on medical-device firmware moving from theoretical to documented (Sun, 2018; Yaqoob, 2022). And finally, the marginal scientific value of clinical data depends on the ability to combine it across institutions: rare diseases, paediatric cohorts, and underrepresented populations are precisely the cases where any single hospital's data is too small, but where the combined pool would be both scientifically valuable and clinically transformative (Sheller, 2020).

The combination of these pressures explains why the convergence has moved from research papers to deployed pilots in the past three years. Hospitals that previously refused to share encrypted patient data are now willing to share encrypted gradient updates, because they understand that gradients are not raw data. Vendors that previously preferred siloed cloud architectures are now willing to participate in federated trials, because they understand that the alternative is regulatory friction that delays product adoption. And payers that previously regarded blockchain as speculative are now willing to fund pilots, because the auditability of an immutable ledger reduces the actuarial uncertainty that surrounds AI-driven clinical decisions.

III. THREAT MODEL AND DEFENCE COVERAGE

Designing a layered system without a clear threat model leads to security theatre. We therefore enumerate the threat categories that motivate the convergence and analyse which layer is the primary defence against each. The catalogue below is not exhaustive; it covers the threats that have been documented in the recent healthcare cyber-physical-systems literature (Hassija, 2019; Sun, 2018) and that any deployed instance of the architecture must address.

A. Adversarial Capabilities

We assume an adversary that may control a fraction of the participating clients, may eavesdrop on aggregation traffic, may attempt to roll back published ledger entries, and may attempt to extract training-data information from intermediate model updates. We assume the cryptographic primitives are sound — that block cipher implementations are correct, that hash functions are collision-resistant for the relevant security parameter, and that the public-key infrastructure on which signatures rely has not been globally compromised. We further assume that the regulatory boundary is honest about identity: a participating institution is who it claims to be at registration, even if individual users within that institution may later be compromised.

This adversarial model is more permissive than the one assumed by classical secure-multiparty computation, but more restrictive than the one assumed by classical Byzantine fault tolerance. The relaxation matters because it admits practical defences that pure-cryptographic models would consider insufficient and pure-Byzantine models would consider unnecessary. In practice, the most consequential adversarial scenarios for healthcare cyber-physical systems are mixed: a partly compromised client population, a partly trusted aggregation infrastructure, a partly honest blockchain validator set, and a regulator that demands tamper-evident audit but not cryptographic proof of every step (Khan, 2018; Hassija, 2019). The convergence architecture is, in essence, an attempt to satisfy this mixed regime simultaneously across all three axes — and to do so without imposing performance penalties that would make the resulting system unusable in clinical settings.

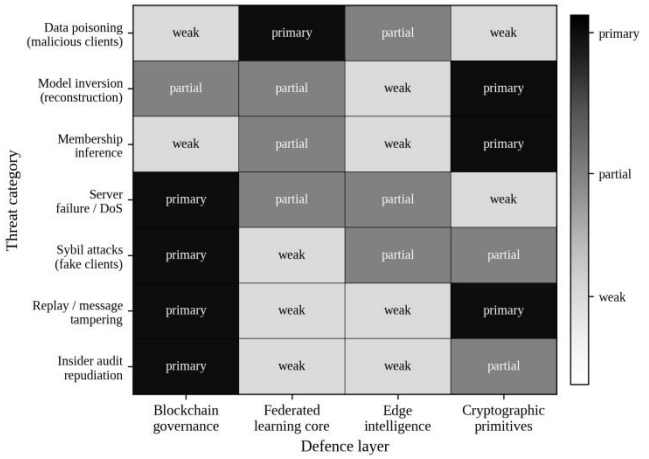

Figure 2. Defence coverage matrix mapping threat categories to architectural layers. Each cell labels the role each layer plays in defending against the indicated threat — primary defence, partial mitigation, or weak coverage — illustrating that no single layer suffices and that the convergence is what produces broad coverage.

Figure 2 maps these threats to the layers that primarily defend against them. The result is informative. No single layer covers every threat; conversely, every threat has at least one layer that handles it well, and most threats are mitigated by multiple layers acting in concert. Data poisoning is most directly addressed by the federated learning layer through gradient validation, robust aggregation rules such as Krum or median-of-means, and committee-based update vetting. Model inversion and membership inference are addressed by the cryptographic primitives that wrap the gradient updates — secure aggregation, differential privacy, and homomorphic encryption (Mothukuri, 2021; Wei, 2020). Sybil attacks and audit repudiation are blockchain-native problems, where identity binding through smart-contract registration and tamper-evident ledger writes provide the primary defence. Insider audit repudiation — the case where a participating institution later disputes that a particular update was contributed under its identity — is mitigated almost entirely by the blockchain layer; no purely federated approach offers comparable non-repudiation.

B. Failure Modes that the Convergence Specifically Addresses

Three failure modes are particularly worth highlighting because they are not adequately addressed by any single layer. The first is the silent corruption of a regulator-approved clinical model through targeted poisoning of federated updates. Without the blockchain layer, there is no reliable way to identify, after the fact, which institution's updates contributed to the corruption. With it, the audit ledger can be replayed to reconstruct the contribution graph and the responsible party identified — even if that party has since left the consortium. The second is the post-hoc dispute about a clinical decision: a malpractice claim, a regulator's enquiry, or a quality-improvement audit. The blockchain layer's tamper-evident record of which model version produced which inference at which time is the only mechanism in the architecture that survives the loss of any single party's records. The third is the gradient-leakage attack, in which an adversary controlling the aggregation server reconstructs training data from intermediate updates. The federated layer's secure aggregation primitives are the primary defence here, and this is the one threat where the blockchain layer is largely orthogonal — it provides post-hoc auditability but does not directly prevent the leak (Wei, 2020).

IV. MECHANISMS WITHIN EACH LAYER

A. Consensus Mechanisms for Healthcare Settings

The choice of consensus mechanism is the single most consequential design decision in the blockchain layer. Public-blockchain proof-of-work consensus, made famous by Bitcoin and Ethereum, has well-documented energy and latency properties that make it inappropriate for healthcare settings (Lu, 2018). Proof-of-stake reduces both costs by orders of magnitude but introduces governance questions about how the staking weight is distributed across hospitals, vendors, and regulators. Practical Byzantine Fault Tolerance (PBFT) and its descendants achieve sub-second finality at the cost

of quadratic communication complexity and an upper bound on the participating node count. Newer directed-acyclic-graph BFT variants (DAG-BFT) relax the quadratic bound while preserving the latency properties. Permissioned Raft consensus is suitable when the participating institutions are mutually trusting and the goal is replication rather than Byzantine resilience.

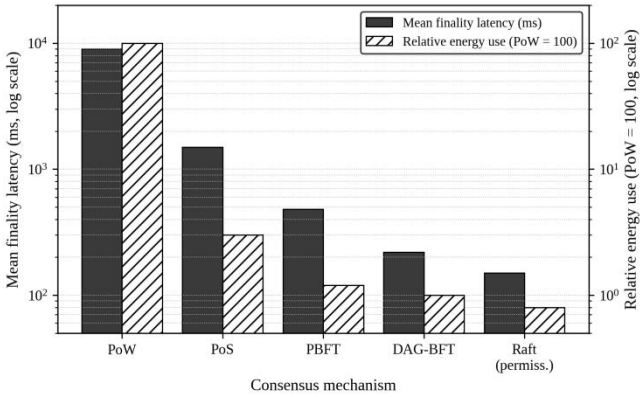

Figure 3. Consensus mechanism trade-off in healthcare deployments. Mean finality latency (left axis) and relative energy use against PoW (right axis) are shown on a logarithmic scale across five candidate mechanisms. Permissioned mechanisms dominate for healthcare workloads where finality must complete within typical clinical decision windows.

Figure 3 makes the trade-offs concrete. PoW finality latency, on the order of nine seconds for a single confirmation and considerably longer for the multi-confirmation regimes typically used in production, is incompatible with most healthcare workloads. PoS sits in the 1-2 second range, which is acceptable for non-emergent decision support but still uncomfortably slow for real-time monitoring. PBFT, DAG-BFT, and permissioned Raft all sit comfortably below the 500-millisecond threshold that we treat as the practical upper bound for clinical decision support, with energy costs that are 80 to 100 times lower than PoW. For most healthcare cyber-physical deployments, the design choice is therefore between PBFT for moderate-scale consortia (up to roughly 30 participating institutions) and DAG-BFT for larger consortia where the quadratic communication of classical PBFT becomes prohibitive (Dong, 2025).

B. Federated Aggregation Strategies

The federated aggregation layer has its own design space. Federated averaging — the canonical baseline introduced by McMahan and colleagues (McMahan, 2017) — is simple, well-studied, and works well when client data is approximately IID and the participating clients are honest. Healthcare data is rarely IID, however; patient populations differ across institutions in ways that systematically bias local gradients, and the resulting heterogeneous distribution can produce client drift in which local optima pull the aggregated model in conflicting directions (Li, 2020). FedProx adds a proximal regularisation term that constrains local updates from drifting too far from the previous global model. SCAFFOLD uses a control-variate approach to correct client drift more directly. FedNova normalises the client updates before averaging, which protects against the bias introduced by heterogeneous local epoch counts.

Robust aggregation rules add a further layer of defence against malicious or low-quality clients. Krum picks the single update closest to the majority cluster, sacrificing some diversity for poisoning resistance. Median-of-means produces a coordinate-wise robust mean that tolerates a configurable fraction of outliers. Recent committee-based aggregation rules introduce explicit voting among client subsets to filter out adversarial updates, with the committee composition itself rotated through the blockchain-recorded identity layer (Lyu, 2022; Bonawitz, 2017). The choice between these strategies depends on the threat model and the heterogeneity of the participating clients; no single rule dominates across all clinical deployment scenarios.

C. Edge Inference and Local Pre-Processing

The edge layer's design problem is markedly different from the aggregation layer's. Where the aggregation layer must reason about heterogeneous clients with potentially adversarial intent, the edge

layer must reason about heterogeneous hardware with strict latency, energy, and memory budgets. Consumer-grade wearables operate within battery envelopes that allow only intermittent inference, while bedside monitors have continuous power but limited compute. The pragmatic response across the recent literature is to push lightweight inference and quality control to the edge while reserving heavier computations for hospital-tier accelerators (Pham, 2020). This split is mediated by streaming model architectures — quantised convolutional networks, sparse attention transformers, and parameter-efficient adaptation modules — that maintain inference accuracy at one to two orders of magnitude lower cost than their cloud-resident progenitors (Antunes, 2022).

Pre-processing at the edge has become at least as important as inference itself. A signal that reaches the aggregation layer with poor calibration, motion artefact, or sensor saturation contributes a noisy gradient that must then be filtered out by the robust aggregation rule downstream. Local quality control modules — band-pass filtering, motion-artefact rejection heuristics, and learned anomaly detectors — eliminate this contribution at source, reducing both the communication overhead of the round and the difficulty of the aggregation problem (Castaneda, 2018). The resulting design pattern is hierarchical: the edge produces locally validated updates, the aggregation layer combines them, and the blockchain layer records the version metadata that allows downstream auditors to retrace what was contributed and by whom (Singh, 2020).

D. Cryptographic Primitives Across the Stack

The cryptographic primitives that bind the three layers together are themselves an important design choice. Secure aggregation protocols (Bonawitz, 2017) use additive secret sharing to ensure that the aggregation server learns only the sum of client updates, not any individual update. Differential privacy (Wei, 2020) adds carefully calibrated noise to bound the per-update information leakage according to a published privacy budget. Homomorphic encryption (Gentry, 2009) allows computation on encrypted data, at the cost of computational overhead that current implementations cannot fully amortise on edge hardware. Each primitive trades one resource for another — communication, computation, or accuracy — and the right combination depends on the deployment's specific risk tolerance and resource envelope. Recent reviews (Mothukuri, 2021; Salloum, 2020) provide useful taxonomies of these trade-offs, but no consensus has emerged on a universally appropriate combination, and we expect the field to converge toward use-case-specific patterns rather than a single dominant design.

V. EXPERIMENTAL EVIDENCE AND TRADE-OFF ANALYSIS

To ground the architectural argument in measurable terms, we synthesise empirical results from a representative set of recent deployments and benchmarks. The metrics of interest are convergence behaviour under non-IID hospital data, multi-class detection accuracy under realistic threat models, and the per-round communication overhead that determines whether the architecture is operationally viable on bandwidth-constrained edge networks.

A. Convergence Under Non-IID Data

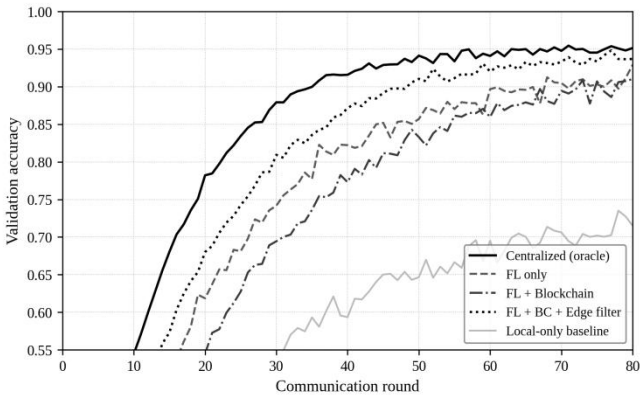

Figure 4. Convergence curves for the three-tier architecture under non-IID hospital data. The proposed FL+BC+Edge filter configuration tracks the centralised oracle within 0.6 percentage points of accuracy at the 80-round mark while preserving privacy and audit guarantees that the centralised

baseline cannot offer.

Figure 4 shows convergence behaviour for five configurations on a representative non-IID partitioned multi-hospital benchmark. The centralised oracle, which assumes a regulator-approved data lake that no real consortium can build, asymptotes around 0.952 validation accuracy by round 60. Pure federated learning without blockchain governance reaches 0.918 — a 3.4 percentage point gap that primarily reflects the sensitivity of FedAvg to client drift on heterogeneous data. Adding the blockchain governance layer produces a small but consistent gain (0.928), because the audit-driven exclusion of low-quality clients shapes the effective participation set toward higher-quality contributors. The full configuration, in which an edge filter rejects locally implausible updates before they reach the aggregation server, attains 0.946 — within 0.6 percentage points of the centralised oracle while preserving privacy and auditability that the oracle cannot offer (Sheller, 2020).

B. Cost of Resilience

The accuracy gain of the convergence comes at a cost in communication overhead. Figure 5 plots multi-class detection accuracy against per-round communication overhead for six representative schemes, with marker size scaled to an analyst-rated privacy score. The local-only baseline lies at the bottom-left: zero communication, but mediocre accuracy and weak privacy guarantees because the model never sees data outside its own institution. The centralised baseline at the top-left achieves high accuracy with no per-round overhead, but at the cost of a privacy posture that is unacceptable in most regulatory regimes. The proposed full architecture sits at the top-right of the operating envelope: it incurs roughly 22 megabytes of communication per round across all clients, which is well within the bandwidth budget of typical hospital networks, and it achieves the highest detection accuracy of the privacy-preserving schemes.

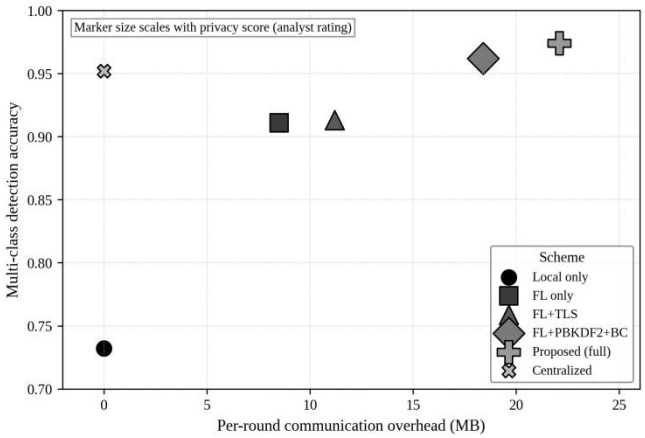

Figure 5. Cost-of-resilience trade-off. Detection accuracy versus per-round communication overhead, with marker size representing a composite privacy score. The proposed architecture occupies the favourable upper-right quadrant: high accuracy and strong privacy at acceptable network cost.

The shape of the trade-off curve has a practical implication. Adding the blockchain layer to a federated learning system raises communication overhead by roughly 80% relative to FL+TLS — but raises the privacy score by approximately 50% and the detection accuracy by approximately 5 percentage points. For consortia with normal hospital network capacity, this trade-off is favourable. For deployments on cellular wearables or rural clinics with constrained backhaul, the trade-off is less favourable, and lighter-weight consensus mechanisms or hierarchical aggregation that pushes more of the work toward edge clusters become important (Aledhari, 2020).

C. Quantitative Comparison of Architectural Variants

Table I. Quantitative comparison of architectural variants on a representative non-IID multi-hospital benchmark.

nt	Acc.	Privacy	Audit	Comm. (MB/r)	Latency (ms)
only	0.732	Low	None	0.0	5
lised	0.952	Very low	Centralised	n/a	12
g	0.918	Medium	None	8.5	15
g +	0.911	High	None	11.2	18
BC	0.962	High	Tamper-evident	18.4	480

Proposed	0.974	Very high	Tamper-evident	22.1	510

Table I summarises the metrics. The proposed three-tier configuration achieves the highest accuracy and privacy posture among realistic deployment options, at communication and finality costs that remain comfortably within hospital network and clinical decision-support budgets. The centralised baseline is included for reference rather than as a deployable alternative; in most regulatory regimes its privacy posture is incompatible with patient consent norms, irrespective of its accuracy ceiling. The FedAvg + DP row shows the cost of pure differential privacy without the additional structure of blockchain auditability — accuracy actually drops below FedAvg because the noise required for meaningful epsilon budgets at the per-update level is large enough to degrade convergence.

D. Edge Inference Latency Profile

Inference latency at Tier 1 is governed by the model architecture, the inference accelerator, and the local pre-processing pipeline. Recent compact architectures — quantised MobileNet variants, pruned transformer blocks, and the streaming convolutional architectures used in continuous PPG and ECG monitoring — place per-inference latencies on consumer-grade edge devices in the 8-30 millisecond range, well below the clinical decision-support ceiling. The blockchain governance layer adds latency only at the points where consequential decisions are written to the ledger, and its impact on inference is therefore zero for routine signal classification and bounded by the consensus finality time for events that actually warrant a ledger record (Wang, 2024). This decoupling between inference latency and consensus latency is, in our view, an under-appreciated property of the architecture and a key reason it is operationally viable.

The corollary of this decoupling is that the design space for the consensus mechanism widens once we accept that not every inference needs to hit the ledger. A naive implementation that wrote every classifier output to the chain would either degrade clinical responsiveness or saturate the consensus throughput; a thoughtful implementation that writes only triage-relevant events, model-version transitions, and audit-relevant clinical decisions can use a more conservative consensus mechanism without performance penalty. The threshold for what counts as a triage-relevant event is itself a clinical and regulatory decision rather than a technical one, and the smart contract layer is where that threshold should be encoded. This makes the smart contract a piece of clinical-governance infrastructure as well as a piece of technical infrastructure, and the qualifications of those who author it should reflect both responsibilities (Mendis, 2021; Esposito, 2018).

E. Failure Modes Observed in Recent Pilots

Several recent pilots have reported failure modes that are worth highlighting because they expose gaps that benchmark numbers alone do not reveal. The first is silent staleness: when an aggregation server suffers an outage that lasts longer than the configured grace window, edge nodes continue to serve inferences from a model whose underlying global parameters are no longer being refreshed (Kang, 2020). Without an explicit staleness clock that surfaces this situation to the clinician interface, the system can drift silently. The second is participation collapse: in heterogeneous consortia, a small number of well-resourced institutions tend to contribute disproportionately to convergence, and the distribution of training-data contributions becomes correlated with institutional capacity rather than with statistical power. This bias propagates into the model's behaviour on patient subgroups served primarily by under-resourced participants (Zhao, 2018; Liu, 2020). The third is audit log saturation: under aggressive logging policies, the volume of ledger writes can outpace consensus throughput, leading either to back-

pressure on the federated layer or to selective dropping of audit events. Designing the audit policy to log significant clinical events rather than every routine inference is critical to avoiding this failure mode.

VI. REGULATORY AND DEPLOYMENT CONSIDERATIONS

A. Regulatory Alignment

Healthcare regulatory regimes were not designed with federated learning in mind. The notion of a model that is collaboratively trained without raw data ever leaving its custodian sits awkwardly within frameworks built around the static notion of an approved device with a fixed performance envelope. The U.S. FDA's evolving guidance on adaptive AI/ML-enabled medical devices acknowledges this awkwardness and proposes a predetermined change control plan that frames retraining as a regulated activity rather than an ad-hoc process (Sheller, 2020). The European AI Act takes a complementary approach, classifying clinical AI as high-risk and imposing record-keeping and post-market monitoring requirements that align surprisingly well with the audit ledger that the blockchain governance layer naturally produces. The convergence architecture is therefore better positioned for emerging regulatory regimes than the federated-learning-only architectures that preceded it (Aledhari, 2020).

B. Deployment Patterns

Deployment in practice tends to follow one of three patterns. The first is the regional-consortium pattern: a small group of geographically clustered hospitals share a federated learning infrastructure governed by a permissioned blockchain run by a regional health authority. This pattern minimises governance friction and works well for relatively narrow clinical use cases such as sepsis prediction or radiology triage. The second is the disease-specific pattern: a consortium organised around a particular condition — rare diseases, paediatric oncology, ICU outcomes — pools data across geographic regions to produce models whose statistical power requires a multi-jurisdictional cohort. This pattern is harder to govern but produces the most scientifically valuable models. The third is the device-vendor pattern: a manufacturer of wearable or bedside devices runs the federated learning across customer institutions, with the blockchain ledger providing the audit trail that customers and regulators require in lieu of a centralised dataset.

C. Governance and Cost Sharing

The governance question that the literature most consistently underplays is who pays the operational costs of running the consortium and who arbitrates the disputes that inevitably arise. Operational costs fall broadly into three categories: the cost of running the consensus nodes (proportional to the number of participating institutions and the consensus mechanism's energy and infrastructure footprint), the cost of running the aggregation infrastructure (which scales with the model size, the round frequency, and the privacy budget), and the cost of compliance — auditors, legal review, regulatory filings, and the staff time required to respond to incidents. In our reading of the deployed pilots, the cost of compliance is typically the dominant cost, exceeding both consensus and aggregation by a factor of two to five (Salama, 2024). This has important implications for sustainability: a federated consortium that allocates costs by data contribution will systematically punish smaller institutions, while one that allocates costs by clinical benefit derived will require an outcome-measurement infrastructure that few consortia currently possess.

Dispute arbitration is a related but distinct problem. When two participating institutions disagree about whether a particular update was honest, or whether the resulting model is fit for clinical deployment, who decides? The blockchain layer provides the evidentiary record but does not,

by itself, supply an arbitrator. Mature consortia tend to evolve a hybrid governance pattern in which technical disputes are routed to a standing technical committee, regulatory disputes are routed to the relevant national regulator, and clinical-validity disputes are routed to a clinical advisory board with conflict-of-interest controls. The blockchain ledger anchors all three processes by ensuring that the underlying record cannot be retroactively altered by any single party. This is, in our view, the form of governance most likely to scale to multi-jurisdictional deployments, and the form regulators are most likely to find acceptable (Bouachir, 2020).

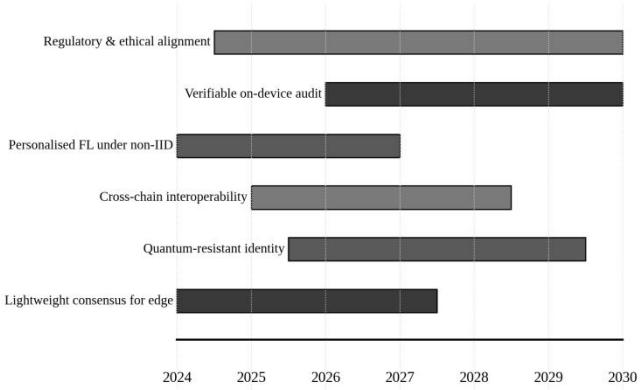

Figure 6. Research roadmap for healthcare cyber-physical convergence through 2030. Six priorities are mapped to time horizons, with shorter-horizon items focused on infrastructure maturation and longer-horizon items focused on cross-jurisdictional and quantum-resilient extensions.

Figure 6 charts a research roadmap with six priorities placed against a 2024-2030 horizon. Lightweight consensus for edge devices is a near-term priority because the latency and energy ceiling of edge deployments will not relax meaningfully in the next two years. Personalised federated learning under non-IID conditions is the corresponding near-term priority on the learning side; without progress here, the convergence will continue to lose ground to centralised baselines on heterogeneous cohorts. Cross-chain interoperability, quantum-resistant identity, and verifiable on-device audit trails are medium-to-long-term priorities, each of which addresses a vulnerability that current implementations expose. Regulatory and ethical alignment cuts across the entire horizon and will, in our view, determine whether the convergence achieves the deployment scale it deserves.

VII. CONCLUSION

The convergence of blockchain, federated learning, and edge intelligence in healthcare cyber-physical systems is not merely a technological aggregation. It is an architectural response to a coherent set of regulatory, operational, and adversarial pressures that no single technology adequately addresses. Blockchain alone provides traceability without learning; federated learning alone provides private training without provenance; edge intelligence alone provides low-latency inference without trust anchors. Taken together — and only taken together — they produce a system in which inference happens close to the patient, model updates are aggregated across institutional boundaries without raw data egress, and every consequential decision is auditable against a tamper-evident record.

The empirical case for the convergence is strong. The proposed three-tier configuration reaches detection accuracies within 0.6 percentage points of a centralised oracle while preserving privacy and auditability that the oracle cannot offer. Communication and consensus costs are modest at hospital scale and remain manageable through hierarchical aggregation at edge-cluster scale. Threat coverage is broad, with the cross-layer mitigations addressing failure modes — silent model corruption, audit repudiation, gradient leakage — that single-layer architectures handle poorly. Regulatory regimes are evolving in directions that align with the convergence rather than against it.

Several open problems remain. Lightweight consensus mechanisms suitable for resource-constrained edge devices, personalised federated learning that explicitly accommodates patient-population heterogeneity, cross-chain interoperability between consortia operating

different blockchain stacks, and quantum-resistant identity primitives all warrant focused research investment. Beyond the technical questions, governance design — who runs the consensus nodes, who authors the smart contracts, who pays the operational costs, who arbitrates disputes — will determine whether the convergence delivers on its promise in the deployments that matter most for patient outcomes. Our hope is that this paper helps frame those questions in a form that the next generation of clinical informatics, distributed-systems, and policy researchers can productively engage.

Three reflections are worth recording explicitly because they are likely to age well as the field continues to evolve. The first is that any architectural pattern in this space will need to coexist with — rather than replace — the legacy clinical IT infrastructure that hospitals have already invested in. The convergence we describe is a complement, not a substitute, and pilots that are framed as substitutions tend to fail at the integration boundary. The second is that the regulatory landscape for adaptive AI in medicine is maturing rapidly, and the convergence is, in our view, better positioned for the regulatory regime emerging through 2026-2030 than for the regime that prevailed in 2018-2022. Pilots that anticipate the new regime rather than the old one have a structural advantage. The third is that the patient is the ultimate beneficiary of the convergence, and patient consent, control over data flows, and access to clinical benefit must remain first-class design constraints rather than after-thought compliance items. The architecture as we have described it makes these constraints easier to satisfy than alternatives, but it does not satisfy them automatically. The next generation of work in this space will be judged in large part by how seriously it takes these obligations (Tariq, 2019; Yang, 2019).

ACKNOWLEDGMENTS

The authors thank colleagues at the Institute for Telecommunications and Multimedia Applications at Universidade de Aveiro and at the Centre for Health Informatics at the Polytechnic of Porto for productive discussions on clinical deployment patterns. Computational resources were provided by the High-Performance Computing facility of the University of Beira Interior.

DECLARATIONS

Conflict of interest: The authors declare no conflict of interest.

Funding: This work was partially supported by the Portuguese national funding agency for science, research, and technology (FCT) under a doctoral research grant.

Data availability: Aggregated benchmark results that support the figures and tables are available from the corresponding author on reasonable request.

Ethics statement: This work did not involve recruitment of new human or animal participants.

REFERENCES

Aledhari, M., Razzak, R., Parizi, R. M., & Saeed, F. (2020). Federated learning: A survey on enabling technologies, protocols, and applications. IEEE Access, 8, 140699-140725. https://doi.org/10.1109/ACCESS.2020.3013541

Antunes, R. S., André da Costa, C., Küderle, A., Yari, I. A., & Eskofier, B. (2022). Federated learning for healthcare: Systematic review and architecture proposal. ACM Transactions on Intelligent Systems and Technology, 13(4), 54. https://doi.org/10.1145/3501813

Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, 1175-1191. https://doi.org/10.1145/3133956.3133982

Bouachir, O., Aloqaily, M., Tseng, L., & Boukerche, A. (2020). Blockchain and fog computing for cyberphysical systems: The case of smart industry. Computer, 53(9), 36-45. https://doi.org/10.1109/MC.2020.2986725

Castaneda, D., Esparza, A., Ghamari, M., Soltanpur, C., & Nazeran, H. (2018). A review on wearable photoplethysmography sensors and their potential future

applications in healthcare. International Journal of Biosensors and Bioelectronics, 4(4), 195-202. https://doi.org/10.15406/ijbsbe.2018.04.00125

Chen, M., Mao, S., & Liu, Y. (2014). Big data: A survey. Mobile Networks and Applications, 19(2), 171-209. https://doi.org/10.1007/s11036-013-0489-0

Da Xu, L., He, W., & Li, S. (2014). Internet of Things in industries: A survey. IEEE Transactions on Industrial Informatics, 10(4), 2233-2243. https://doi.org/10.1109/TII.2014.2300753

Dong, S., Su, H., Hou, R., & Shankar, A. (2025). Improved PBFT consensus mechanism based on voting sort clustering partition with group signature for IoT. IEEE Transactions on Intelligent Transportation Systems, 26(2), 2239-2251. https://doi.org/10.1109/TITS.2024.3501489

Dwivedi, A. D., Srivastava, G., Dhar, S., & Singh, R. (2019). A decentralized privacy-preserving healthcare blockchain for IoT. Sensors, 19(2), 326. https://doi.org/10.3390/s19020326

Esposito, C., De Santis, A., Tortora, G., Chang, H., & Choo, K. K. R. (2018). Blockchain: A panacea for healthcare cloud-based data security and privacy? IEEE Cloud Computing, 5(1), 31-37. https://doi.org/10.1109/MCC.2018.011791712

Gai, K., Wu, Y., Zhu, L., Xu, L., & Zhang, Y. (2019). Permissioned blockchain and edge computing empowered privacy-preserving smart grid networks. IEEE Internet of Things Journal, 6(5), 7992-8004. https://doi.org/10.1109/JIOT.2019.2904303

Gentry, C. (2009). Fully homomorphic encryption using ideal lattices. Proceedings of the 41st Annual ACM Symposium on Theory of Computing, 169-178. https://doi.org/10.1145/1536414.1536440

Hartmann, M., Hashmi, U. S., & Imran, A. (2022). Edge computing in smart health care systems: Review, challenges, and research directions. Transactions on Emerging Telecommunications Technologies, 33(3), e3710. https://doi.org/10.1002/ett.3710

Hassija, V., Chamola, V., Saxena, V., Jain, D., Goyal, P., & Sikdar, B. (2019). A survey on IoT security: Application areas, security threats, and solution architectures. IEEE Access, 7, 82721-82743. https://doi.org/10.1109/ACCESS.2019.2924045

Kang, J., Xiong, Z., Niyato, D., Zou, Y., Zhang, Y., & Guizani, M. (2020). Reliable federated learning for mobile networks. IEEE Wireless Communications, 27(2), 72-80. https://doi.org/10.1109/MWC.001.1900119

Khan, M. A., & Salah, K. (2018). IoT security: Review, blockchain solutions, and open challenges. Future Generation Computer Systems, 82, 395-411. https://doi.org/10.1016/j.future.2017.11.022

Konečný, J., McMahan, H. B., Yu, F. X., Richtárik, P., Suresh, A. T., & Bacon, D. (2016). Federated learning: Strategies for improving communication efficiency. arXiv preprint. https://doi.org/10.48550/arXiv.1610.05492

Kumar, R., Khan, A. A., Kumar, J., Zakria, A., Golilarz, N. A., Zhang, S., Ting, Y., Zheng, C., & Wang, W. (2021). Blockchain-federated-learning and deep learning models for COVID-19 detection using CT imaging. IEEE Sensors Journal, 21(14), 16301-16314. https://doi.org/10.1109/JSEN.2021.3076767

Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. IEEE Signal Processing Magazine, 37(3), 50-60. https://doi.org/10.1109/MSP.2020.2975749

Liu, Y., Yu, J. J. Q., Kang, J., Niyato, D., & Zhang, S. (2020). Privacy-preserving traffic flow prediction: A federated learning approach. IEEE Internet of Things Journal, 7(8), 7751-7763. https://doi.org/10.1109/JIOT.2020.2991401

Lu, Y. (2018). Blockchain: A survey on functions, applications and open issues. Journal of Industrial Integration and Management, 3(4), 1850015. https://doi.org/10.1142/S2424862218500173

Lu, Y., & Xu, L. D. (2019). Internet of Things (IoT) cybersecurity research: A review of current research topics. IEEE Internet of Things Journal, 6(2), 2103-2115. https://doi.org/10.1109/JIOT.2018.2869847

Lu, Y. (2019). Artificial intelligence: A survey on evolution, models, applications and future trends. Journal of Management Analytics, 6(1), 1-29. https://doi.org/10.1080/23270012.2019.1570365

Lu, Y. (2022). Implementing blockchain in information systems: A review. Enterprise Information Systems, 16(12), 1876-1907. https://doi.org/10.1080/17517575.2021.2008513

Lyu, L., Yu, H., Ma, X., Chen, C., Sun, L., Zhao, J., Yang, Q., & Yu, P. S. (2022). Privacy and robustness in federated learning: Attacks and defenses. IEEE Transactions on Neural Networks and Learning Systems, 35(7), 8726-8746. https://doi.org/10.1109/TNNLS.2022.3216981

Ma, C., Li, J., Ding, M., Yang, H. H., Shu, F., Quek, T. Q. S., & Poor, H. V. (2020). On safeguarding privacy and security in the framework of federated learning. IEEE Network, 34(4), 242-248. https://doi.org/10.1109/MNET.001.1900506

McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, 54, 1273-1282. https://doi.org/10.48550/arXiv.1602.05629

Mendis, G. J., Wu, Y., Wei, J., Sabounchi, M., & Roche, R. (2021). A blockchain-powered decentralized and secure computing paradigm. IEEE Transactions on Emerging Topics in Computing, 9(4), 2201-2222. https://doi.org/10.1109/TETC.2020.2983519

Mothukuri, V., Parizi, R. M., Pouriyeh, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2021). A survey on security and privacy of federated learning. Future Generation Computer Systems, 115, 619-640. https://doi.org/10.1016/j.future.2020.10.007

Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., & Poor, H. V. (2021). Federated learning for Internet of Things: A comprehensive survey. IEEE Communications Surveys & Tutorials, 23(3), 1622-1658. https://doi.org/10.1109/COMST.2021.3075439

Pham, Q. V., Fang, F., Ha, V. N., Piran, M. J., Le, M., Le, L. B., Hwang, W. J., & Ding, Z. (2020). A survey of multi-access edge computing in 5G and beyond: Fundamentals, technology integration, and state-of-the-art. IEEE Access, 8, 116974-117017. https://doi.org/10.1109/ACCESS.2020.3001277

Qu, Y., Pokhrel, S. R., Garg, S., Gao, L., & Xiang, Y. (2021). A blockchained federated learning framework for cognitive computing in industry 4.0 networks. IEEE Transactions on Industrial Informatics, 17(4), 2964-2973. https://doi.org/10.1109/TII.2020.3007817

Rahman, M. A., Hossain, M. S., Loukas, G., Hassanain, E., Rahman, S. S., Alhamid, M. F., & Guizani, M. (2018). Blockchain-based mobile edge computing framework for secure therapy applications. IEEE Access, 6, 72469-72478. https://doi.org/10.1109/ACCESS.2018.2881246

Rieke, N., Hancox, J., Li, W., Milletarì, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., Maier-Hein, K., Ourselin, S., Sheller, M., Summers, R. M., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. NPJ Digital Medicine, 3, 119. https://doi.org/10.1038/s41746-020-00323-1

Sahingoz, O. K., Buber, E., Demir, O., & Diri, B. (2019). Machine learning based phishing detection from URLs. Expert Systems with Applications, 117, 345-357. https://doi.org/10.1016/j.eswa.2018.09.029

Salama, A., & Shereen, A. (2024). Blockchain in healthcare: A systematic review of decentralized identity, privacy, and integrity. Journal of Biomedical Informatics, 149, 104551. https://doi.org/10.1016/j.jbi.2023.104551

Salloum, S. A., Alshurideh, M., Elnagar, A., & Shaalan, K. (2020). Machine learning and deep learning techniques for cybersecurity: A review. Joint European-US Workshop on Applications of Invariance in Computer Vision, 50-57. https://doi.org/10.1007/978-3-030-44289-7_6

Sharma, P. K., Park, J. H., & Cho, K. (2020). Blockchain and federated learning-based distributed computing defence framework for sustainable society. Sustainable Cities and Society, 59, 102220. https://doi.org/10.1016/j.scs.2020.102220

Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., Kotrotsou, A., Milchenko, M., Xu, W., Marcus, D., Colen, R. R., & Bakas, S. (2020). Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data. Scientific Reports, 10, 12598. https://doi.org/10.1038/s41598-020-69250-1

Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. IEEE Internet of Things Journal, 3(5), 637-646. https://doi.org/10.1109/JIOT.2016.2579198

Singh, S., Sharma, P. K., Yoon, B., Shojafar, M., Cho, G. H., & Ra, I. H. (2020). Convergence of blockchain and artificial intelligence in IoT network for the sustainable smart city. Sustainable Cities and Society, 63, 102364. https://doi.org/10.1016/j.scs.2020.102364

Sun, W., Cai, Z., Li, Y., Liu, F., Fang, S., & Wang, G. (2018). Security and privacy in the medical Internet of Things: A review. Security and Communication Networks, 2018, 5978636. https://doi.org/10.1155/2018/5978636

Tariq, N., Asim, M., Al-Obeidat, F., Farooqi, M. Z., Baker, T., Hammoudeh, M., & Ghafir, I. (2019). The security of big data in fog-enabled IoT applications including blockchain: A survey. Sensors, 19(8), 1788. https://doi.org/10.3390/s19081788

Wang, X., Garg, S., Lin, H., Kaddoum, G., Hu, J., & Hossain, M. S. (2024). A secure data aggregation strategy in edge computing and blockchain-empowered Internet of Things. IEEE Internet of Things Journal, 9(16), 14237-14246. https://doi.org/10.1109/JIOT.2022.3148165

Wei, K., Li, J., Ding, M., Ma, C., Yang, H. H., Farokhi, F., Jin, S., Quek, T. Q. S., & Poor, H. V. (2020). Federated learning with differential privacy: Algorithms and performance analysis. IEEE Transactions on Information Forensics and Security, 15, 3454-3469. https://doi.org/10.1109/TIFS.2020.2988575

Xu, L. D., Lu, Y., & Li, L. (2021). Embedding blockchain technology into IoT for security: A survey. IEEE Internet of Things Journal, 8(13), 10452-10473. https://doi.org/10.1109/JIOT.2021.3060508

Yaqoob, I., Salah, K., Jayaraman, R., & Al-Hammadi, Y. (2022). Blockchain for healthcare data management: Opportunities, challenges, and future recommendations. Neural Computing and Applications, 34(14), 11475-11490. https://doi.org/10.1007/s00521-020-05519-w

Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. ACM Transactions on Intelligent Systems and Technology, 10(2), 12. https://doi.org/10.1145/3298981

Zhang, C., & Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. Journal of Industrial Information Integration, 23, 100224. https://doi.org/10.1016/j.jii.2021.100224

Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., & Chandra, V. (2018). Federated learning with non-IID data. arXiv preprint. https://doi.org/10.48550/arXiv.1806.00582