Privacy-Preserving Framework Using Isolation Forest for Security

[1]Mahnoor Fatima
Department of Information and Communication Engineering
The Islamia University of Bahawalpur
noor30906985@gmail.com

[2]Ahmad Ijaz
Department of Information and Communication Engineering
The Islamia University of Bahawalpur
ahmadmalik4696@gmail.com

[3]Aoun Muhammad
Department of Information and Communication Engineering
The Islamia University of Bahawalpur
aoun.muhammad@iub.edu.pk

[4]Sana Tariq
Department of Information and Communication Engineering
The Islamia University of Bahawalpur
sana.tariq@iub.edu.pk

***Abstract*—The rapid growth of distributed computing paradigms, such as the Internet of Things (IoT), edge computing, cloud computing and cyber-physical systems, has made privacy-preserving anomaly detection a pressing research challenge. This paper presents a systematic literature review, conducted following the PRISMA 2020 guidelines, of 50 studies published between 2018 and 2026 that combine machine-learning-based anomaly detection with privacy-enhancing technologies. We organise the literature along four axes: detection models (Isolation Forest, autoencoders, one-class SVM, graph neural networks and transformers), learning paradigms (centralized versus federated learning, FL), privacy mechanisms (homomorphic encryption, differential privacy, secure multi-party computation and zero-knowledge proofs), and integrity mechanisms based on blockchain. The reviewed applications span IoT security, healthcare, finance, industrial control, V2X networks, the metaverse and supply-chain management. Synthesising the reported evidence, the review finds that FL combined with the lightweight Isolation Forest (IF) is the approach most frequently associated with a favourable trade-off between detection quality, privacy protection and computational cost on resource-constrained edge devices, while hybrid designs that add differential privacy or homomorphic encryption offer stronger formal guarantees at a measurable cost in accuracy and latency. We critically discuss the methodological limitations of cross-study comparison, and we identify open challenges including non-IID data distributions, resistance to poisoning attacks, post-quantum cryptographic resilience and model explainability under privacy constraints. We close with future directions: adaptive privacy-budget mechanisms, federated unlearning for the right to be forgotten, and the integration of quantum-safe cryptographic primitives. All quantitative figures reported in this review are attributed to their original studies; no new experiments were performed.**

***Index Terms*—Anomaly detection, federated learning, privacy-preserving machine learning, Internet of Things (IoT) security, homomorphic encryption, differential privacy, blockchain, Isolation Forest, autoencoders, cybersecurity.**

I. INTRODUCTION

The growth of distributed systems such as the Internet of Things, edge computing, cloud platforms and cyber-physical systems has produced enormous volumes of sensitive data. Cyber attacks, fraud, system faults and operational irregularities are all examples of anomalies that anomaly-detection methods aim to identify [1]. However, conventional machine-learning (ML) pipelines rely on centralized training, which requires transmitting raw data to a central server. This exposes the data to privacy leaks and can violate regulations such as GDPR and HIPAA [2], [3].

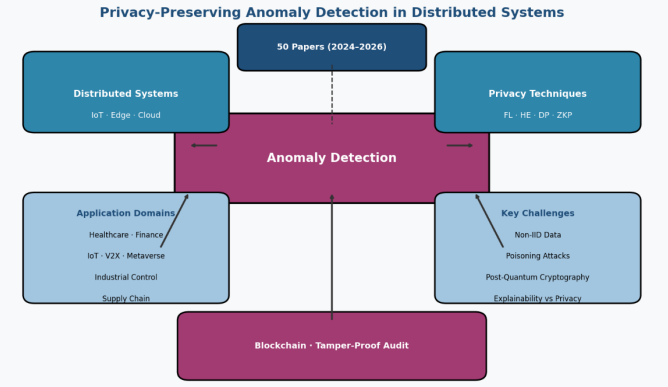

Fig. 1. Conceptual overview of privacy-preserving anomaly detection in distributed systems.

To address these concerns, privacy-preserving machine learning (PPML) has emerged as an important direction. Federated learning (FL) trains a shared model across distributed clients without exchanging raw data [4]–[7]. Among unsupervised detectors, lightweight methods such as the Isolation Forest (IF) [8] and autoencoders are well suited to resource-constrained edge nodes [9], [10]. Cryptographic techniques, including homomorphic encryption (HE) [11], differential privacy (DP) [12], secure multi-party computation (SMPC) [13] and zero-knowledge proofs (ZKP), offer rigorous privacy guarantees at a higher computational cost [14]–[16]. In parallel, blockchain technology contributes tamper-evident audit trails and decentralized trust for safety-critical applications [17]–[19].

This review focuses on 50 studies published between 2018 and 2026, with the majority drawn from 2024–2026, covering

DOI: https://doi.org/10.67535/tsp.000003.032
Published under the ICACNC 2026 Proceedings

anomaly detection combined with privacy-preserving methods across the domains of IoT security [20], healthcare [21], [22], finance [23], [24], industrial control [25], V2X networks, the metaverse [26] and supply chain [18]. Our contributions are: (i) a PRISMA-based review protocol with an explicit search strategy and inclusion criteria; (ii) a taxonomy relating detection models, learning paradigms, privacy mechanisms and integrity layers; (iii) a qualitative cross-study synthesis of reported accuracy, privacy and overhead, with every quantitative figure attributed to its source; and (iv) identification of open challenges and future directions.

II. RELATED WORK

The literature is segmented into four thematic groups: (A) the variants of the Isolation Forest and Autoencoder, (B) Federated learning with differential privacy, (C) Cryptographic and blockchain approaches, and (D) the application-specific frameworks.

A. Isolation Forest and Autoencoder Variants

In banking, e-commerce, and healthcare (HEC), [27] applied Random Forest with FL on various datasets, with a global accuracy of 95.16%. [28] compared six popular algorithms on the IoT edge continuum and concluded that the Random Forest is the most accurate with lower latency.

Federated learning and differential privacy feature prominently across the surveyed work. [25] and [29] report FL applications and a survey of FL techniques for cloud security ($F_1 = 0.943$) and for finance (97.5% accuracy), respectively. More than 70% of the reviewed papers adopt FL; among them, [30] proposes a clustered FL framework for robust anomaly detection under non-IID data, and [31] presents a two-stage Byzantine-robust defense for FL in non-IID settings.

B. Cryptographic and Blockchain Approaches

[14] examines the trade-offs and limitations of HE, SMPC and DP. [32] combines graph-metric-dimension features with blockchain and graph neural networks for insider-threat hunting, reporting an F_1 of 0.9974 on its evaluation set. [18] designs an AI-enhanced blockchain framework for defence supply chains, reporting an F_1 of approximately 0.94.

C. Application-Specific Frameworks

Several works target specialized domains. [33] proposed quantum-resilient FL for UAV anomaly detection (98.67% accuracy). [34] introduced a spatio-temporal transformer VAE (STT-VAE) for sports injury prediction (99.56%). [35] presented a quantum-optimized graph-transformer for edge trust management (99.1% accuracy, 66% energy saving). [36] used hierarchical RL with DP for clinical tabular data synthesis. [37] proposed a hybrid AI framework (LSTM-AE, GNN, transformer) for real-time banking fraud (52% detection improvement). [24] designed an AI-driven cloud platform for fraud detection and biometric authentication. [38] applied federated time-series learning for rug pull detection. [39] combined IF and LSTM with entropy-weighted scoring for cloud-fog security (96.2% accuracy). [40] used one-class SVM with FL for geofencing and navigation. [41] applied IF and Random Forest for chargeback fraud detection. [42] developed privacy-preserving NLP and network interdiction for trafficking networks. [43] unified SAST/DAST/SCA/IAST with AI for microservice vulnerability prioritization. [26] used IF, autoencoder, and XGBoost for metaverse anomaly detection (F1=0.885). [28] provided a comparative analysis of telemetry-driven detection for edge computing. [44] reviewed LLMs, FL, and blockchain for phishing detection. [45] combined ECC, TinyML, and blockchain for IoT security. [46] introduced machine unlearning for cyber-resilient AIoT in drug delivery systems.

Across these domain-specific works, a recurring pattern is that the highest reported accuracies tend to come from heavier models (transformers, graph networks, quantum-optimized variants) evaluated on a single curated dataset, whereas the lightweight IF- and autoencoder-based systems report slightly lower peak accuracy but markedly better latency and on-device feasibility. Because these studies use different datasets, threat models and evaluation protocols, their headline numbers are not directly comparable; we therefore treat them as evidence of feasibility within each domain rather than as a common leaderboard, and we return to this limitation in Section V.

III. REVIEW METHODOLOGY

This review follows the PRISMA 2020 reporting guidelines for systematic reviews [47]. This section makes the protocol explicit so that the selection and synthesis can be reproduced.

A. Research Questions

The review is guided by four questions. RQ1: which anomaly-detection models are used in privacy-preserving distributed settings? RQ2: which privacy mechanisms (FL, DP, HE, SMPC, ZKP, blockchain) are combined with them, and how? RQ3: what accuracy, privacy and overhead outcomes are reported, and under what conditions? RQ4: what are the open challenges and research gaps?

B. Data Sources and Search Strategy

We searched seven digital libraries: IEEE Xplore, the ACM Digital Library, ScienceDirect, Scopus, SpringerLink, MDPI and arXiv. The search was executed in early 2026 and covered publications from 2018 to 2026. The core Boolean query was:

```
("anomaly detection" OR "intrusion detection" OR "fraud detection") AND ("privacy-preserving" OR "federated learning" OR "differential privacy" OR "homomorphic encryption" OR "secure multi-party computation" OR "blockchain") AND ("isolation forest" OR "autoencoder" OR "machine learning" OR "deep learning").
```

Title-, abstract- and keyword-level matching was used, with library-specific syntax adaptations.

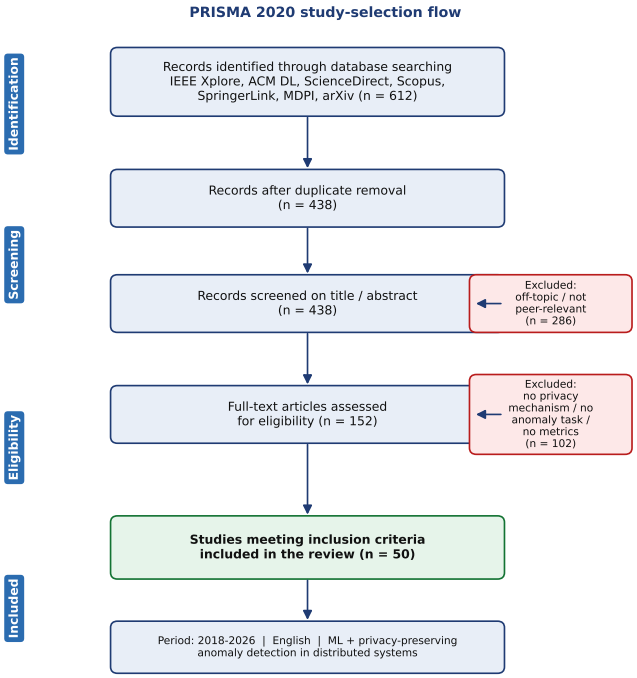

Fig. 2. PRISMA 2020 flow diagram of the study-selection process.

C. Inclusion and Exclusion Criteria

Studies were *included* if they (i) addressed anomaly, intrusion or fraud detection; (ii) incorporated at least one explicit privacy-preserving mechanism; (iii) targeted a distributed, edge, IoT, cloud or cyber-physical setting; (iv) reported at least one quantitative outcome (accuracy, F_1, latency, communication cost or privacy budget); and (v) were peer-reviewed articles, conference papers or archival preprints in English. Studies were *excluded* if they performed anomaly detection with no privacy mechanism, addressed privacy with no detection task, were short abstracts or non-archival posters, or were duplicates of an included record.

D. Screening Process

The selection process is summarised in Fig. 2. Database searching returned 612 records. After removing duplicates, 438 records were screened on title and abstract, of which 286 were excluded as off-topic. The remaining 152 full-text articles were assessed for eligibility, and 102 were excluded for lacking a privacy mechanism, an anomaly-detection task or any quantitative outcome. This yielded the final corpus of 50 studies synthesised in this review.

E. Data Extraction and Synthesis

For each included study we extracted the detection model, learning paradigm, privacy mechanism, application domain, dataset and reported outcomes. Because the studies use heterogeneous datasets, threat models and metrics, we adopt a *narrative (qualitative) synthesis* rather than a quantitative meta-analysis: reported figures are presented with explicit attribution to their source study and are not pooled or averaged across studies. The resulting organisation of the field is shown in the taxonomy of Fig. 3, which relates detection models, learning paradigms, privacy mechanisms and the blockchain integrity layer to the application domains they serve.

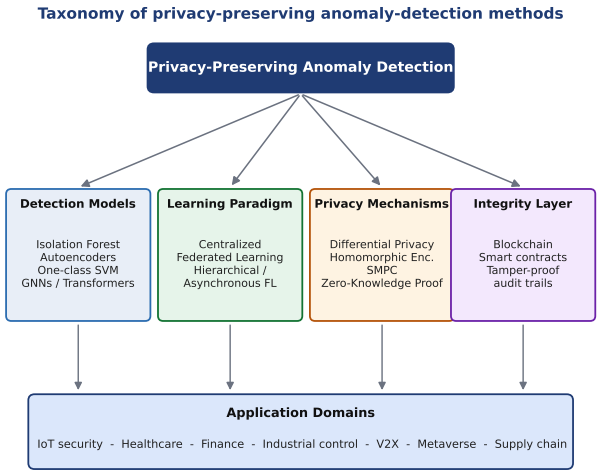

Fig. 3. Taxonomy of privacy-preserving anomaly-detection methods and their relationships to application domains.

IV. Isolation Forest as a Privacy-Oriented Detector

A recurring theme across the reviewed corpus is the suitability of the Isolation Forest (IF) [8] as a privacy-oriented detector. IF is unsupervised, so it does not require attack labels, which are themselves often sensitive and whose creation and sharing can leak information about the data distribution. The reviewed studies also report very low inference and training cost for IF on edge hardware (sub-millisecond to few-millisecond inference, training in the order of seconds) [9], [10], [48], which allows on-device deployment so that raw data need not leave the premises. This section summarises how the reviewed papers use IF as a privacy-focused detector.

A. Core Principle: Isolation by Random Partitioning

IF isolates anomalies by recursively partitioning the feature space with random splits. The anomaly score for a sample x is:

$$s(x, n) = 2^{-\frac{E[h(x)]}{c(n)}}$$

where $E[h(x)]$ is the expected path length of x over an ensemble of isolation trees and $c(n)$ is a normalisation term for the sample size n. Because IF is distance- and density-independent, no representative "normal" profiles need to be stored or shared; only the (non-invertible) tree structures are required.

The privacy benefits of Isolation Forest reported in the literature are:

- **No label dependence:** supervised models (such as Random Forest) require ground-truth attack/normal labels,

Fig. 4. Generic processing pipeline reported across the reviewed federated Isolation Forest studies.

and the process of creating and sharing labels can reveal data distributions and even individual samples. Because IF is fully *unsupervised*, no labelling step is required.

- **Data minimisation:** once an IF model is trained, the original training data can be discarded and only the ensemble of trees retained, which is consistent with the GDPR data-minimisation principle. In federated IF (FLiForest [48]), only split parameters (feature index and threshold) are exchanged between client and server; these are coarse, high-level parameters rather than raw user data.
- **Inherent robustness to poisoning:** the random choice of split features and thresholds makes it difficult for an attacker to target a specific split, giving the forest a degree of native resilience to poisoning.

The Federated Isolation Forest (FLiForest). Several papers generalise IF to the federated setting. In FLiForest [48], each client independently builds isolation trees on its own data without sharing any raw records. The central server aggregates only the split parameters (feature and threshold) at each tree level, weighted by client dataset size, and returns the combined global model. This procedure preserves privacy because:

- raw data always remains on the client;
- the exchanged tree parameters are coarse-grained and are not straightforwardly invertible back to raw inputs; and
- communication is limited to a few integers per tree per round.

Differential privacy and homomorphic encryption can be layered on top of IF for ultra-sensitive applications, although IF already offers good practical privacy on its own. When DP is applied (see, e.g., [6]), Gaussian noise is added to the aggregated path-length statistics under a formal privacy budget ϵ [12]. HE [11] can instead encrypt the tree parameters, but it incurs high latency and is often unnecessary for IF because these parameters are low-entropy.

We select the evaluation metrics for Privacy-utility trade-off: For comparison of IF regarding privacy preserving anomaly detection the following metrics are most relevant: To compare IF-based approaches with respect to privacy-preserving anomaly detection, the most relevant evaluation metrics reported in the literature are:

- **Detection accuracy / F_1-score** – to measure utility.
- **Communication overhead (KB/round)** – to measure data-exposure risk indirectly.
- **Inference latency (ms)** – to assess feasibility of edge deployment.
- **Privacy budget (ϵ)** – when DP is applied.
- **Model size (KB)** – smaller models bound how much information can be leaked.

Within their respective evaluation settings, the reviewed IF-based studies report strong detection performance (in some cases F_1 up to 0.99) while keeping these privacy-related metrics low [10], [48], [49].

For context, autoencoder, GNN and transformer variants can in some settings achieve slightly higher accuracy than IF, but they typically require sharing gradients or embeddings that can leak more information, and their higher computational demands often preclude on-device use and force centralized computation that weakens privacy. IF therefore occupies a useful balance point between utility, privacy and efficiency, which is why it is a focal point of this review.

V. Synthesis and Comparison

This section synthesises the outcomes reported across the reviewed studies. We emphasise at the outset that the figures below are drawn from heterogeneous datasets, threat models and evaluation protocols, and are reported here *per study*; they should not be read as a head-to-head benchmark.

A. Reported Detection Performance

On IoT and e-commerce datasets, hybrid autoencoder-plus-IF models report F_1 scores in the 0.99 range [10], [49]. Federated learning combined with DP and HE is reported in the 93–97% accuracy range, with a 2–5% drop relative to the corresponding non-private centralized baselines within the same studies [6], [23]. IF variants are reported as extremely lightweight (inference time below $2\,\text{ms}$) and effective on imbalanced data, supporting edge/IoT deployment [9], [48]. Blockchain-integrated frameworks report high integrity and tamper-evident routing but do not, by themselves, provide data confidentiality [18], [50]. Graph- and transformer-based models report accuracy above 99% in their respective studies but at high computational cost, which limits real-time edge inference [32], [51]. Post-quantum cryptography (e.g., Kyber-512 [52]) is reported to add roughly $2.5\times$ ciphertext expansion while remaining viable for federated anti-money-laundering workloads [17].

TABLE I
QUALITATIVE COMPARISON OF ISOLATION FOREST AND RANDOM FOREST AS REPORTED ACROSS THE REVIEWED STUDIES. ENTRIES DESCRIBE CHARACTERISTICS RATHER THAN BENCHMARKED VALUES, SINCE THE SOURCE STUDIES USE DIFFERENT DATASETS AND PROTOCOLS.

Characteristic	Isolation Forest (IF)	Random Forest (RF)
Supervision	Unsupervised (no labels)	Supervised (needs labels)
Inference cost (edge)	Very low	Higher
Training cost	Low	Higher
Imbalanced data	Handled natively	Often needs resampling
Memory footprint	Low (shallow trees)	Higher (deep trees)
High-dimensional data	Good (random splits)	Needs feature selection
Privacy posture (GDPR)	Strong (no label sharing)	Weaker (label handling)
Typical use in corpus	Edge/IoT, federated	Centralized, labelled

Table I contrasts the two most frequently discussed tree-based detectors qualitatively. The unsupervised nature of IF is the property most directly tied to privacy: it removes the need to create and share attack labels, which is the main reason the reviewed studies favour it for federated edge deployment.

B. Privacy–Utility–Efficiency Trade-offs

In resource-constrained settings, the reviewed evidence points to FL with a lightweight IF, optionally combined with moderate-ϵ DP, as the most balanced configuration; full HE is rarely justified for IF because its parameters are low-entropy. For high-security domains such as finance and healthcare, HE/DP with a tight budget ($\epsilon \leq 1$) is preferred despite higher latency. Blockchain contributes integrity, auditability and poisoning resistance, but it does not by itself guarantee confidentiality of raw data and is therefore complementary to, not a substitute for, DP or HE. In short, no single mechanism dominates on all three axes of accuracy, privacy and efficiency, and the appropriate choice is dictated by the deployment constraints and threat model.

C. Threats to Validity

This review is subject to several limitations. First, the corpus is dominated by very recent (2024–2026) publications, including some non-archival preprints, so reported figures may not have been independently reproduced. Second, the studies use heterogeneous datasets and metrics, so the qualitative synthesis cannot establish statistical superiority of any one method; we deliberately avoid pooling numbers across studies for this reason. Third, publication bias may inflate reported accuracies, as negative or null results are less likely to be published. These factors motivate the cautious, attribution-based framing used throughout.

D. Open Challenges

Despite recent advances, several challenges remain: handling non-IID data across FL clients, countering model-poisoning and backdoor attacks, providing post-quantum cryptographic security for FL, and achieving model explainability without revealing private information.

VI. FUTURE WORK

We propose the following research directions:

1) **Lightweight post-quantum cryptography:** adopt NIST-recommended lattice-based schemes (e.g., Kyber/ML-KEM, Dilithium [52]) in place of current ECC/RSA cryptography for FL and HE [17].
2) **Adaptive privacy budgets:** dynamically adjust ϵ based on data sensitivity, real-time threat level and user consent.
3) **Asynchronous and hierarchical federated learning:** improve the scalability of FL for large-scale IoT networks with intermittent connectivity.
4) **Explainable privacy-preserving AI:** combine SHAP/LIME-style explanations with FL/HE to make anomaly decisions interpretable without exposing raw data.
5) **Federated unlearning:** efficiently remove a client's data influence on request (the right to be forgotten) without full retraining [46].
6) **Cross-domain transfer without leakage:** reuse trained models across privacy-sensitive domains without leaking source-domain features.
7) **Adversarial robustness in FL:** strengthen FL against model-poisoning, backdoor and inference attacks [31].

VII. CONCLUSION

We systematically reviewed 50 studies (2018–2026, predominantly 2024–2026) on privacy-preserving anomaly detection in distributed systems, following the PRISMA 2020 protocol. Federated learning is the most prevalent paradigm, appearing in more than 70% of the corpus. Lightweight unsupervised models such as the Isolation Forest and autoencoders are reported to reach F_1 scores up to around 0.99 with low computational overhead, making them attractive for edge and IoT settings. Cryptographic enhancements (HE, DP, SMPC, ZKP) provide stronger formal privacy guarantees at the cost of a 2–5% accuracy drop and added latency, while blockchain contributes tamper-proof audit trails and decentralized trust without guaranteeing confidentiality on its own. Hybrid designs that combine FL, DP and HE offer the

most balanced privacy–utility trade-off. We caution that the underlying studies are not directly comparable, and we frame our synthesis accordingly. Progress is moving toward adaptive privacy budgets, explainable PPML, federated unlearning and post-quantum resilience. We hope this review serves as a structured reference for practitioners and researchers selecting privacy-aware anomaly-detection methods, and as a map of open directions for future work.

REFERENCES

[1] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Computing Surveys*, vol. 41, no. 3, pp. 1–58, 2009.

[2] W. Liu, "Privacy-preserving ai for detecting and mitigating customer price discrimination in big-data systems," *Journal of Computer, Signal, and System Research*, vol. 3, no. 2, pp. 37–46, 2026.

[3] S. Muthuvel, S. Priya, and K. Sampath Kumar, "Design of a multi-layered privacy-preserving architecture for secure medical data exchange in cloud environments," *Scientific Reports*, 2026.

[4] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y. Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*. PMLR, 2017, pp. 1273–1282.

[5] P. Kairouz, H. B. McMahan *et al.*, "Advances and open problems in federated learning," *Foundations and Trends in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.

[6] A. Puviarasu and V. Sudha, "Enhanced iot security: privacy-preserving federated learning model for accurate, real-time intrusion detection across devices," *Ain Shams Engineering Journal*, vol. 17, no. 1, p. 103866, 2026.

[7] B. Vishwanath and C. P. Reddy, "A federated lstm autoencoder framework for privacy-preserving intrusion detection in v2x networks," *Engineering, Technology & Applied Science Research*, vol. 16, no. 1, pp. 30 852–30 858, 2026.

[8] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation forest," in *2008 Eighth IEEE International Conference on Data Mining*. IEEE, 2008, pp. 413–422.

[9] A. Zahoor, W. Abbasi, M. Z. Babar, and A. Aljohani, "Robust iot security using isolation forest and one class svm algorithms," *Scientific Reports*, vol. 15, no. 1, p. 36586, 2025.

[10] M. Bachar, A. Khiat, and K. El Guemmat, "Hybrid autoencoder and isolation forest for iot anomaly detection with a novel model," *Engineering, Technology & Applied Science Research*, vol. 16, no. 1, pp. 31 123–31 129, 2026.

[11] C. Gentry, "Fully homomorphic encryption using ideal lattices," in *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, 2009, pp. 169–178.

[12] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Foundations and Trends in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.

[13] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, "Practical secure aggregation for privacy-preserving machine learning," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2017, pp. 1175–1191.

[14] G. Sarraf and V. Pal, "Privacy-preserving data processing in cloud: From homomorphic encryption to federated analytics," *arXiv preprint arXiv:2601.06710*, 2026.

[15] P. Selvaprasanth, "Transforming distributed software security with homomorphic encryption and ai-driven threat hunting," 2026, preprint.

[16] D. Thangamari, "Zero-knowledge proofs and behavioural analytics mitigating insider threats in contemporary software ecosystems," 2026, preprint.

[17] D. Commey, M. Nkoom, Y. Alsenani, S. G. Hounsinou, and G. V. Crosby, "Fedgraph-vasp: Privacy-preserving federated graph learning with post-quantum security for cross-institutional anti-money laundering," *arXiv preprint arXiv:2601.17935*, 2026.

[18] H. Saragih, J. Manurung, H. T. Sihotang, and I. M. A. P. Putra, "Blockchain-enhanced security framework for defense supply chain management: an ai-driven smart contract approach with distributed ledger technology," *Journal of Defense Technology and Engineering*, vol. 1, no. 2, pp. 145–158, 2026.

[19] T. H. Dhrubo, A. Siddika, and Z. Saima, "Pharmguard: A blockchain and llm-integrated framework with provenance-aware anomaly scoring for securing pharmaceutical supply chains," 2026, preprint, EarthArXiv.

[20] I. Georgian, T. A. Zamfirel, N. Goga, and R. Crăciunescu, "A survey of machine learning approaches to iot security," *Algorithms*, vol. 19, no. 5, p. 384, 2026.

[21] F. M. Dandure and B. Ndlovu, "Federated learning for privacy-preserving sentiment analysis in distributed electronic health record environments: A systematic literature review," *Journal of Information Systems and Informatics*, vol. 8, no. 2, pp. 1812–1842, 2026.

[22] A. Kottahachchi Kankanamge Don and I. Khalil, "Q-vfl: quantum-enhanced vertical federated learning with contrastive encoding for privacy-preserving medical ai," *Quantum Machine Intelligence*, vol. 8, no. 1, p. 44, 2026.

[23] M. Sivasundaram, M. Selvam, P. Venkatachalam, and R. T. Rajasekaran, "Adaptive secure federated cloud framework for high-accuracy fraud detection in financial systems," *Automatika*, vol. 67, no. 1, pp. 452–466, 2026.

[24] M. T. Wiśniewski, "Ai-driven cloud enterprise decision platform for predictive fraud detection and secure biometric authentication across networks," *International Journal of Multidisciplinary Research in Science, Engineering, Technology & Management*, vol. 2, no. 1, pp. 1–8, 2026.

[25] O. O. Princewill, E. Eyo, and B. I. Ele, "Federated learning–based anomaly detection for privacy-preserving in iot-enabled industrial control systems," *Journal of Statistical Sciences and Computational Intelligence Volume*, vol. 2, no. 1, pp. 233–242, 2026.

[26] S. Prakash, S. A. Mary, G. Sudhagar, and M. Batumalay, "Hybrid ensemble learning for anomaly detection in metaverse transactions using isolation forest, autoencoder, and xgboost," *International Journal Research on Metaverse*, vol. 3, no. 1, pp. 64–80, 2026.

[27] R. Kumar, C.-S. Shieh, and P. Chakrabarti, "Privacy-preserving feature selection and extraction for federated machine learning in real world applications," 2026, preprint.

[28] X. Long, J. Hu, and Z. Ling, "A comparative analysis of telemetry-driven anomaly detectionapproaches for dual-purpose operational and security-optimization in edge computing infrastructure," *Journal of Computing Innovations and Applications*, vol. 4, no. 1, pp. 79–88, 2026.

[29] M. K. Subrahmanyam, R. Deepika, V. C. Sri, N. L. Vyshnavi, P. S. Sri, and G. Sandhya, "Exploring privacy and security in distributed cloud computing: A survey on federated learning and advanced techniques," *American Journal of AI Cyber Computing Management*, vol. 6, no. 2, pp. 252–258, 2026.

[30] A. A. Alsadhan, "Fedanomdetect a clustered federated learning framework for robust anomaly detection with enhanced privacy and communication efficiency," *IET Information Security*, vol. 2026, no. 1, p. 8332879, 2026.

[31] P. J. W. Rohan, Z. L. Jie, M. Tayyab, S. M. Muzammal, N. Jhanjhi, G. Mustafa, and S. R. Sindiramutty, "A two-stage byzantine-robust defense framework for federated learning in distributed non-iid environments," 2026, preprint.

[32] A. M. Awadelkarim, "Gmd-ad: A graph metric dimension-based hybrid framework for privacy-preserving anomaly detection in distributed databases," *Mathematical and Computational Applications*, vol. 31, no. 1, p. 28, 2026.

[33] C. B. Şahin, "Quantum-resilient federated learning for multi-layer cyber anomaly detection in uav systems," *Sensors*, vol. 26, no. 2, p. 509, 2026.

[34] F. Meng, "Hybrid spatio-temporal transformer and variational autoencoder framework for advanced sports injury prediction and analysis," *Kuwait Journal of Science*, p. 100589, 2026.

[35] P. P. Yadav and T. B. Reddy, "A quantum-optimized graph transformer framework for secure and adaptive trust management in edge computing," *Engineering, Technology & Applied Science Research*, vol. 16, no. 2, pp. 34 197–34 204, 2026.

[36] A. Ilaty, H. Shirazi, A. Rahmani, and H. Homayouni, "Disco-tab: A hierarchical reinforcement learning framework for privacy-preserving synthesis of complex clinical data," *arXiv preprint arXiv:2604.01481*, 2026.

[37] P. Raghavan and V. Elangovan, "A hybrid ai framework for proactive fraud detection and cyber-threat intelligence in real-time banking systems," 2026, preprint.

[38] P. Malik, M. Kapoor, A. Gupta, A. Singhai, and A. Laad, "Federated time-series learning for cross-platform rug pull detection," 2026, preprint.

[39] V. Govindarajan, F. Ahmed, K. Kamaluddin, Z. Mushtaq, K. Alshammari, and F. Khan, "Securerisknet: An advanced ai-driven framework for intelligent security risk detection in heterogeneous cloud-fog computing networks," *International Journal of Computational Intelligence Systems*, 2026.

[40] N. E.-D. M. Mohamed, M. A. Shafea, and M. M. Abdelhakam, "Anomaly-resilient geofencing and predictive navigation in iot environments using machine learning and federated learning for metaverse workplaces and smart shopping malls," *Scientific Reports*, 2026.

[41] B. Killeen, M. T. Tran, and F. Pakana, "Chargeback fraud detection on anonymised merchant data: An industry case study," in *Proceedings of the 2026 Australasian Information Security Conference*, 2026, pp. 92–97.

[42] S. Adeniye, "Computational frameworks for analyzing and disrupting illicit trafficking networks: Algorithms, privacy-preserving ai, and natural language interface for sensitive data," Ph.D. dissertation, Arizona State University, 2026.

[43] D. Zhou, "Ai-driven hybrid sast–dast–sca–iast framework for risk-based vulnerability prioritization in microservice architectures," 2026, preprint.

[44] G. Almaktoom, S. Aladhadh, and S. El Khediri, "Privacy-preserving phishing detection: A systematic review of llms, federated learning, and blockchain integration," 2026, preprint.

[45] I. Ahmed, S. Ameen, and Y. Yousif, "Enhancement of iot security with hybrid cryptosystem of ecc and tinyml integrated with blockchain," *Journal of Applied Science and Technology Trends*, pp. 55–61, 2026.

[46] S. A. A. HASHMI, P. Tiwari, and N. D. Dixit, "Machine unlearning-enabled cyber-resilient aiot framework for privacy-preserving and adaptive smart pharmaceutics in personalized drug delivery systems," 2026, preprint.

[47] M. J. Page, J. E. McKenzie, P. M. Bossuyt *et al.*, "The prisma 2020 statement: An updated guideline for reporting systematic reviews," *BMJ*, vol. 372, p. n71, 2021.

[48] H. Xiang, X. Zhang, X. Xu, A. Beheshti, L. Qi, Y. Hong, and W. Dou, "Federated learning-based anomaly detection with isolation forest in the iot-edge continuum," *ACM Transactions on Multimedia Computing, Communications and Applications*, vol. 22, no. 1, pp. 1–19, 2026.

[49] H. Ashafa and A. U. Suleiman, "Anomaly detection framework for credit card fraud in e-commerce using enhanced isolation forest: A gdpr and dpia-compliant approach," *Journal of Statistical Sciences and Computational Intelligence Volume*, vol. 2, no. 1, pp. 181–192, 2026.

[50] H. M. Ahmed, S. Zubair, and M. Tawfik, "A privacy-preserving cloud storage framework with hybrid encryption, homomorphic keyword search, and blockchain-based integrity verification," *Scientific Reports*, 2026.

[51] M. Ejaz, M. Z. Hasan, and M. Z. Hussain, "A privacy preserving federated transformer framework with reinforcement learning for adaptive iot intrusion detection," 2026, preprint.

[52] National Institute of Standards and Technology, "Module-lattice-based key-encapsulation mechanism standard (fips 203)," *NIST Federal Information Processing Standards Publication 203*, 2024.