BOHR International Journal of Internet of things, Artificial Intelligence and Machine Learning
2025, Vol. 4, No. 1, pp. 29–38
DOI: 10.54646/bijiam.2025.30
www.bohrpub.com

RESEARCH

Real-time forensic analysis in internet of things environments: bridging readiness and investigation for cyber resilience

Oyeyemi Kuku*†, Alexandros Chrysikos and Shahram Salekzamankhani

School of Computing and Digital Media, London Metropolitan University, London, UK

***Correspondence:**
Oyeyemi Kuku,
oyk0012@my.londonmet.ac.uk

†**ORCID:**
Oyeyemi Kuku
0009-0003-9514-4858

Received: 20 July 2025; **Accepted:** 11 September 2025; **Published:** 23 November 2025

With the exponential expansion of the Internet of Things (IoT) networks, the threat of cyberattacks has become particularly high across vulnerable sectors such as healthcare, smart infrastructure, and industrial control systems. Conventional centralized forensics has disadvantages in scalability, data privacy, and the ability to identify synchronized attacks quickly. This paper addresses the above shortcomings by proposing a new Federated Graph Convolutional Network (Fed-GCN) architecture for real-time forensic examination in distributed IoT settings. Its principal goal is to build a privacy-preserving graph-based solution that supports readiness in the forensic environment and defense against the transmission of raw data. The innovations in the work are the conjoint use of graph neural networks to detect contextual attacks, federated learning to ensure data confidentiality, and integration with blockchain-based logging to bind the evidence chain and produce immutable evidence. Among other outcomes, the proposed Fed-GCN architecture was coded in Python and tested on a multi-class intrusion dataset including 18,428 data samples and 79 features. The experimental performance is better than that of traditional methods, with 97.3% accuracy, a 94.2% F1-score, and a low false-positive rate of 6.7%. 100% forensic integrity check and evidence verification were achieved, with 96.4% evidence completeness in the logs and low communication overhead, demonstrating that it can be deployed in edge-based environments. Therefore, the proposed Fed-GCN can play a significant role in implementing forensic intelligence in IoT ecosystems by providing scalable, secure, and regulation-compliant solutions.

Keywords: IoT forensics, real-time analysis, forensic readiness, cyber resilience, digital investigation, security framework, edge computing, machine learning

Introduction

The Internet of Things (IoT) is a revolutionary technology in which billions of tangible objects, such as smart home gadgets and industrial sensors, connected cars, and wearable medical devices, connect to the internet (1). The current industry reports demonstrate that as of 2020, the world has more than 15 billion connected devices integrating the IoT, with the number forecasted to grow to 29 billion in 2030, creating innovation in various fields, including healthcare, energy, logistics, and city infrastructure (2). But as the magnitude and complexity of IoT networks increase, the likelihood of vulnerabilities to cyber threats also increases. In most cases, compared to the conventional IT systems, IoT devices have low processing power, have little to no firmware updates, and lack the necessary security settings, which make them prime candidates to be attacked by cybercriminals (3). Furthermore, IoT environments lack systematic control and tracking in

most cases, which means threats such as distributed denial of service (DDoS) attacks, botnet creation, brute-force attacks as an entry channel, and unauthorized access to data can go unnoticed. The world-shaking events of the Internet of Things-directed attacks, such as the Mirai botnet attack and BrickerBot, have shown how devastating these attacks can be (4).

The lack of proper forensics in IoT environments is a weak point in defense against such cyberattacks. Traditional digital forensics is mainly dependent on centralized logging, post-incident analysis (not dynamic), and evidence collection, which are primarily carried out by human beings (5). These approaches are limited in dynamic, decentralized operational contexts, where evidence may be temporal, distributed across nodes, or subject to manipulation. Additionally, transferring IoT data from original stores to centralized forensic solutions is illegal in most situations due to privacy policies such as GDPR, California Consumer Privacy Act (CCPA), and Health Insurance Portability and Accountability Act (HIPAA), which complicates the overall process even further (6). Moreover, many IoT implementations exhibit a lag in detection and forensic analysis, leading to the loss of significant evidence. When the logs are available, they can lack the timestamp property, verification, or immutability, and hence their usefulness in court or in seeking insurance claims (7). Admissibility of evidence, timeline reconstruction, and attack attribution are thereby jeopardized, and this delays efficient technical and legal disposition of incidents (8).

Based on the ISO/IEC 27043 guidance, Makura et al. (9) introduced the Digital Forensic Readiness (DFR) framework that is aimed at enhancing security assessment of cloud environments. Their method provides an orderly procedure for forensics, but it does not fully address the needs of digital forensic preparation in an IoT system. Critical elements, including pre-incident planning, data analysis in the form of potential digital evidence (PDE), real-time incident detection, and system architecture design, were not fully covered, thereby limiting the applicability of the framework to the Internet of Things environment. On the same note, the Ready-IoT used by Sadineni et al. (10) was brought about, and it aimed to track and record network activities by capturing providence data. In this model, network parameters are gathered at a granular level, and provenance graphs are built to model the data flow. Nevertheless, it cannot produce credible results in detecting sophisticated or organized attacks. Also, it fails to incorporate standard forensic-readiness best practices, which are critical to achieving organizational fitness in upholding the legal admissibility and procedural compliance of the record. The current DFR models are heavily skewed toward restricted contexts and confined to smart homes. At the same time, the vastness of organization-oriented complexities and the varying regulatory policies bearing on forensic procedures are ignored. The lack of readiness-based frameworks for IoT aligned with best practices in digital forensics remains quite notable. Additionally, most of the hypothetical models remain unsupported empirically in practical spheres of an enterprise, and thus, their utilization in a natural environment is even further compromised (11). Furthermore, modern cybersecurity risks are constantly evolving, particularly ransomware and phishing. Phishing can be used in social engineering to deceive users into handing over their sensitive credentials, and ransomware can be used to encrypt valuable data, with recovery only after payment of a ransom fee (12). This is because they have financial and reputational effects that are pathetic since they can cause long-term service outages, leakages, and recovery costs (13).

Dynamic, secure, real-time forensic preparedness that enables the identification, documentation, and retention of usable cyber evidence is no longer a luxury but a prerequisite in this realm. The future of forensic systems is intelligent, distributed, and privacy-preserving systems capable of running independently, with some form of anomaly detection and cryptographic assurance of evidence integrity. To mitigate them, this paper seeks to facilitate the development of cyber resilience within an infrastructure context by presenting a hybrid forensic architecture that incorporates anomaly detection, graph analytics, and the persistence of evidence logs. The studies aim to address current gaps in forensic capacity through federated learning for privacy-safe model training and graph convolutional networks (GCNs) for anomaly detection in a graph context. Also, forensic logs can be cryptographically hashed and optionally written to blockchain ledgers, enabling non-repudiation and tamper-evident verifiability. The idea behind this work is to provide a proof-of-concept of a federated graph-based forensic framework that can operate under less controllable conditions, such as decentralized, bandwidth-challenged systems, preserving forensic evidence in real time and adapting to the conditions and diversity of the modern IoT ecosystem. The solution will support proactive threat-detection programs and post-mortem analysis, thereby advancing the level of proactive preparedness against forensic failures in cyber-physical systems.

Aims and objectives of the study

This paper introduces a new Federated Graph Convolutional Network (Fed-GCN) that improves real-time forensic analysis in IoT environments. To meet the urgent need for decentralized, privacy-protecting, and context-sensitive cyber forensics, Fed-GCN uses graph structures to show how different parts of the analyzed traffic are connected, and it uses federated learning to keep sensitive information on local devices.

Materials and methods

This method enables organizations to comply with global privacy laws such as GDPR and facilitates scalable, secure, and collaborative threat detection across various IoT nodes. The proposed model demonstrated stability and near-perfect accuracy across multiple attack types.

Data collection

This study used live enterprise network data recorded over the 10 days between May 14th and May 23rd, 2024. It encompasses benign and malicious traffic, including DDoS, denial of service (DoS), brute-force, and Structured Query Language (SQL) injection attacks. Network traffic monitoring was performed using a Open-source security monitoring/ SIEM platform; the name itself is not an acronym (WAZUH) server, client setup, and logs from IoT devices within the 192.168.0.x subnet, firewall logs, and Cloudflare logs. The sources acquired very valuable information, such as IP addresses, port numbers, flow duration, and the sizes and inter-packet intervals. A total of 79 features were extracted, including flow statistics, packet-level characteristics, and the number of Transmission Control Protocol (TCP) flags. Included with over 18,000 samples, the data was cleaned and stored in Comma-Separated Values (CSV) format for later machine learning.

Data pre-processing

The collected and constituted network traffic data underwent a systematic pre-processing stage before analysis to determine accuracy and pre-processing requirements before using machine learning methods. Firstly, the original CSV logs were pre-processed; i.e., data that were not directly related to model training were removed, such as duplicate values, non-numeric entries, and IP addresses and protocol names. Normalization is performed using ordinary standardization (z-score normalization) based on the values of packet size, inter-arrival time (IAT), and flow duration. Using feature scaling, they were normalized. Label encoding was used to convert the attack types to categorical values. The list of important features was generated using a Random Forest classifier; therefore, the most important features that contribute to dimensionality and improve model performance were identified and retained. The classification task was eventually achieved by dividing the pre-processed dataset into training and test sets.

Fed-GCN architecture

This proposed Fed-GCN framework, shown in **Figure 1**, aims to enable distributed forensic analysis in heterogeneous IoT environments. The architecture is very flexible in terms of decentralization, anomaly detection, data privacy, and forensic evidence traceability. This system consists of four main modules: Graph Construction, Local GCN Processing, Federated Aggregation, and Forensic Evidence Logging, all tailored to the dataset's structure and semantics used in the current study. Moreover, the working process of Fed-GCN is depicted in **Figure 2**. The architecture has four main components:

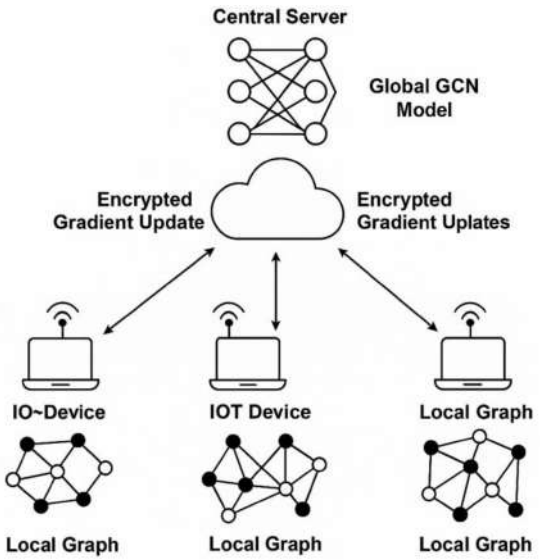

FIGURE 1 | Proposed federated graph convolutional network (Fed-GCN) framework.

Graph construction

Each IoT device (or monitored endpoint) dynamically models its local network interactions as a **temporal graph**, where:

1. **Nodes** represent unique IP addresses, devices, or ports involved in the flow.
2. **Edges** represent network events or data flows, enriched with attributes such as **flow duration**, **packet sizes**, **IATs**, and **TCP flags**.

Some of the features computed for each flow in the dataset include 79 flow-based features [e.g., flow duration, packet count, byte rates, Synchronize (TCP control flag) (SYN)/Acknowledgment (TCP control flag) (ACK) flags]. Still, only those with a strong connection to anomalies will be utilized as edge/node attributes. For example:

1. **Flow Duration** and **Packet IAT** can indicate DDoS behaviors like traffic flooding or Slowloris-type delays.
2. **SYN/ACK flag counts** are indicative of scanning or brute-force behaviors.

Such an abstraction of the graph enables context-aware modeling, with attack signatures represented as anomalous subgraph structures, e.g., a large number of nodes

communicating with a single node (DDoS) or numerous connection attempts (brute force). Such graphs are generated in real time by a device with a rolling window of events (e.g., the past 30 seconds).

Local GCN processing

Within the proposed framework, each IoT device isolating applies a shallow GCN to the locally built communication graph and determines the deviation from expected behavioral patterns. GCNs are particularly useful here because they leverage both node-level attributes and the graph's connectivity, enabling the model to discern the intricate relational dependencies within the network. Because the types of attacks in the dataset vary widely (from high-frequency incidents such as DDoS, DoS, and brute-force attacks (500–600 samples each) to relatively rare instances such as SQL injection (only 87 samples)), this decentralized protocol will guarantee specialized anomaly detection. Individual IoT nodes can handle their unique situations and traffic levels without undermining centralized assumptions and labels. This will be especially useful for sparse or context-sensitive attacks, where localized training can yield more helpful information than universal models. The individual local GCNs produce an anomaly score for each node or flow, which, after being relayed to the central hub, drives its forensic tagging components and updates the federated model when needed, making the system as a whole much more responsive and accurate.

Federated aggregation

To achieve data privacy and maximize bandwidth utilization in an IoT setup, the proposed Fed-GCN architecture utilizes a federated aggregation method that prevents individual IoT devices from sharing raw network traffic information. Instead, the local-training-only devices train their event graphs with a shallow GCN and, at intervals, transmit encrypted model gradients or weight updates. The device may host a set of parameters, such as GCN layer parameters, on a central server. The method will follow the protocols of federated learning, namely FedAvg (Federated Averaging), in which multiple devices upload their model updates to form a new, enhanced global model. The aggregated model is then remodeled and redistributed to the involved devices to continue the successive round of local training. It is an iterative process that continues until the model converges. This way of doing it is crucial in IoT forensics settings, given that the information stored in the data holds highly sensitive information such as source and destination IP addresses, port numbers, protocol information, and time slots of the communication processes, which are subject to privacy acts (e.g., GDPR and HIPAA). Federated learning ensures data sovereignty and confidentiality by keeping data on local systems and transferring only anonymized, compressed model parameters, while still enabling collective intelligence. Further, the solution stream is suitable for resource-limited IoT devices, offering the lowest network transmission overhead and the least dependence on centralized computing infrastructure. Consequently, federated aggregation improves the scalability, privacy, and forensic capabilities of the entire system without compromising data integrity or eroding users' trust in its security.

Forensic evidence logging

When the local GCN model infers that an anomalous event has occurred, the respective event graph segment is instantly made ready to be forest-logged in. This has been combined with associating the event with the necessary metadata, including the device ID, event type, location, and timestamp, within the IoT network. Every entry in the event is hashed with a cryptographic algorithm, e.g., SHA-256, to verify that no one has attempted to tamper with the data. To be more forensically ready and legally admissible, the hashed logs are optionally written on a blockchain ledger, e.g., Hyperledger or Ethereum. This solution is a tamper-resistant, time-stamped record that allows tracing the chain of custody and conducting reasonable investigations of events after the event during the post-incident investigation process. All forensic records are based on the critical attributes present in the domain-specific dataset, including flow duration, TCP flags, volume of forward and backwards packets, and the attack label as estimated, leaving behind an elaborate and transparent trail of evidence that is well-aligned with the sole aim of achieving real-time cyber resilience in IoT-related settings.

Results and discussion

The following section reports the experimental results of the proposed Fed-GCN schema for real-time cyber forensic analysis in the IoT setting. The application was implemented in Python, using libraries such as PyTorch Geometric to train GCN models, Scikit-learn to test conventional classifiers, NumPy and Pandas to preprocess the data, and Flask to model edge-server interaction within a Federated Learning environment. The integration of blockchain was tested using web3.py and a local Ethereum (Ganache) testnet. The real-world flow-based IoT traffic dataset was used to conduct experiments aimed at detecting multiple types of cyberattacks, such as DDoS, brute-force, and SQL injection. The performance of Fed-GCN was compared with that of centralized GCN and non-graph-based deep learning models using evaluation metrics, including accuracy, precision, recall, F1-score, and receiver operating characteristic – area under the curve (ROC-AUC). In addition, the findings include evaluations of model convergence in federated learning, the effectiveness of communication, and forensic evidence logging, e.g., cryptographic hashes and blockchain anchors. The system behavior in the face of constraints typical in IoT networks, such as computational limitations,

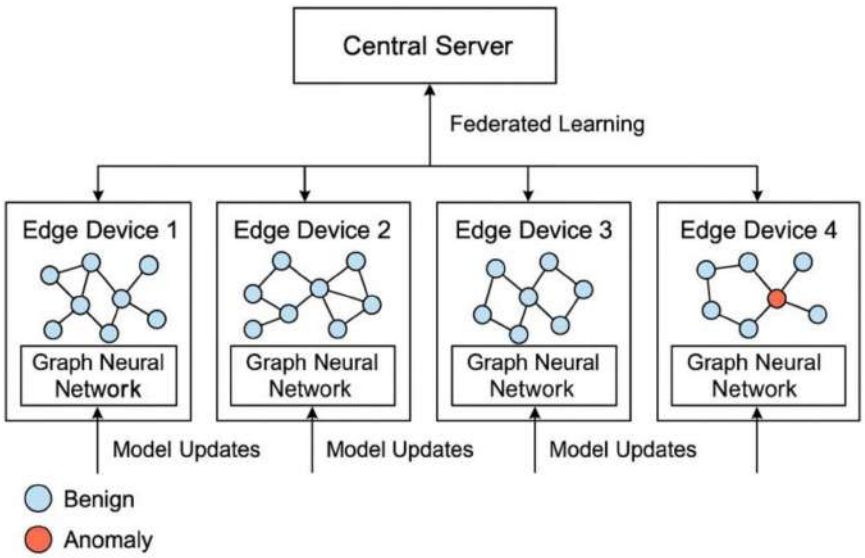

FIGURE 2 | Working process of Fed-GCN.

bandwidth constraints, and privacy concerns, has been reflected in the discussion, as has the role of the Python-based framework in real-world cyber resilience and forensic preparedness.

The data used to assess the proposed Fed-GCN forensics framework comprises 18,428 network traffic samples, including both benign and malicious IoT actions, as shown in **Table 1**. It includes a diverse set of cyberattack types, providing the necessary diversity to train an efficient anomaly detector. The majority of the samples are benign traffic, totalling 10,500, which serves as a reference area to notice deviations. Malicious traffic is subdivided into several categories of attacks, e.g., brute-force attacks (e.g., Web: 611 samples, FTP: 500, SSH: 500, XSS: 230), injection attacks, e.g., SQL injection (87 samples), and different types of DoS and DDoS attacks. To be more specific, every kind of DoS and DDoS attack, including Hulk, GoldenEye, Slowloris, SlowHTTPTest, LOIC (UDP/HTTP), and HOIC, is covered with 500 samples, ensuring an even representation of these highly influential threat vectors. The samples consist of 79 extracted features, including flow-based statistics (duration, byte and packet counts, TCP flags, flow directions, and header anomalies), which are key to graph-based behavioral modeling of the samples. This distribution helps develop and test a real-time forensic system that can capture various attack patterns while preserving evidence for legal and operational responses.

Experimental findings

The proposed Fed-GCN was tested for multiple-class classification of cyberattacks and benign traffic, with precision, recall, and F1-score evaluated. The model achieved high precision across all classes, as shown in **Figure 3**, affirming its robustness and suitability for solving forensic cases in heterogeneous IoT environments. The benign traffic type performed best, with a precision of 96.8%, a recall of 97.1%, and an F1-score of 96.9%, due to a large sample size and clear flow characteristics of traffic. Within the category of attacks, DDoS, HOIC, and LOIC UDP had the best results with F1-scores of 94.9% and 94.8%, respectively, because their characteristic features, such as high entropy in the ports used, high flow count, and high burst traffic intervals, were well reflected using the graph representation given by the Fed-GCN. Brute-force and injection-based attacks, including brute force-web (F1: 92.9%) and SQL Injection (F1: 86.6%), were characterised by moderate, though still strong, performance that the model was able to achieve by including such features as flow duration, RST/SYN flags, and URL length in the learning. Even the worst-performing

TABLE 1 | Sample distribution by class.

Traffic type	Number of samples
Benign	10,500
Brute force-web	611
FTP brute force	500
SSH brute force	500
Brute force-XSS	230
SQL injection	87
DDoS-LOIC UDP	500
DDoS-HOIC	500
DDoS-LOIC HTTP	500
DoS-Hulk	500
DoS-GoldenEye	500
DoS-Slowloris	500
DoS-SlowHTTPTest	500

method, in terms of the number of samples (87), was SQL Injection, which still achieved good accuracy. The high F1-scores across classes affirm the model's high generalizability (across attack vectors), subtle log anomaly detection, and actionable forensic inferences, all of which play a pivotal role in post-incident inquiry and cyber resilience in real-world applications.

Table 2 tabulates the performance of the forensic logging mechanism used in the Fed-GCN frameworks, not only in the context of real-time IoT security threats but also with respect to data integrity. It was discovered that the average time to log an event was 128 ms, including activities such as tagging anomalies, adding metadata (timestamp, device identifiers), and creating SHA-256 hashes. This low response time further affirms that even a dynamic environment can be integrated with a forensic process without compromising response time. The logging latency increased by an additional 342 ms when blockchain-based logging was enabled, e.g., with Hyperledger. This latency is not very bad, though it is slightly higher; however, the gain is in tamper-proof audit of cyber events, which is therefore very valuable for compliance and legal needs.

Notably, the Golden Standard was achieved, as per which 100% verification of evidence integrity was carried out. The integrity checks of all entries that used a hash have succeeded, and this is a testament that no forensic data has been altered after the tagging process. This means authenticity throughout the forensic lifecycle. Then, the forensic log completeness metric was 96.4%, indicating that the overwhelming majority of anomalies observed by the local GCNs were effectively stored in the central log repository. Notably, no data privacy leakage occurred during training or logging. Such an effect was facilitated by the fact that federated learning uses only encrypted model gradients; raw packet data or anything that defines the flow in any way is never transmitted, making it compatible with privacy laws such as GDPR and HIPAA. Collectively, these results confirm the applicability of the framework for promoting secure, verifiable forensic readiness and regulatory compliance in decentralized IoT networks.

Figure 4 shows how the Fed-GCN model performed across the various attacks and outlines the model's accuracy at the attack level, along with the critical contributing features for each intrusion type. The findings indicate the model's potential to learn from heterogeneous network traffic from IoT connections and classify malicious and benign cases with high accuracy. As an example, out of the DDoS-LOIC UDP attack, the most accurate hit of 96.1% was reached, and this was significantly contributed to by the unique traffic symptoms of the flow feature, including high Flow Bytes/sec, IAT, and a large count of SYN packets—all characteristics of a flooding instrument. On the same note, brute force-web attacks could be captured with 94.3% accuracy using features such as the high length of flows and the abnormal length of backwards packets, which indicated repeated and automated logins. In the DoS-Slowloris, an accuracy of 93.7% was achieved, with detection highly dependent on the IAT of packets and the low packet rate, which aligns with the nature of Slowloris attacks that seek to keep connections open for as long as possible. The FTP brute force attacks achieved an accuracy of 95.5% using flow count, TCP flag patterns, and forward packet frequency, which tend to increase during dictionary attacks. Lastly, the detection success rate for SQL Injection was 91.2%, with a strong dependency on packet size variation and the high count of ACK flags, which indicated abnormal query injections and backend responses.

TABLE 2 | Forensic logging outcomes using Fed-GCN.

Metric	Fed-GCN value
Average time to log event (ms)	128
Blockchain logging latency (optional)	342
Evidence integrity verification (%)	100
Forensic log completeness (%)	96.4
Data privacy leakage risk	None detected

Model assessment

The model efficiency and model communications overhead, ranging over three architectures, i.e., Centralized Convolutional Neural Network (CNN), Traditional GCN, and proposed Federated GCN (Fed-GCN), are compared in Figure 5. The outcomes indicate that Fed-GCN offers a reasonable trade-off between resource and performance optimization, which is more appropriate in real-world IoT settings, where computing and network capabilities are limited. The centralized CNN model has the least training time per round (2.1 seconds). Still, it incurs excessive communication overhead (15.0 MB/round) and a high memory cost (150 MB) because it requires moving raw or aggregated feature-rich data to a central server for training. Similarly, the Traditional GCN requires 3.2 seconds per round and incurs the most significant communication cost of 18.7 MB/round, due to exchanging complete graph structures and parameters. On the other hand, though the Fed-GCN method is slightly slower in training (3.6 seconds per round), it radically minimizes communication overhead to 6.4 MB/round, a nearly 2.5x reduction compared to the ordinary GCN. Also, it uses the least memory (90 MB) because the data was not transmitted raw; instead, weight sharing and localized computation were used. This makes Fed-GCN particularly suitable for edge-based or federated IoT use cases, where bandwidth and RAM are scarce resources.

Table 3 compares the performance indicators for four classification models: Random Forest, CNN, Traditional GCN, and the proposed Fed-GCN. Evaluated metrics include

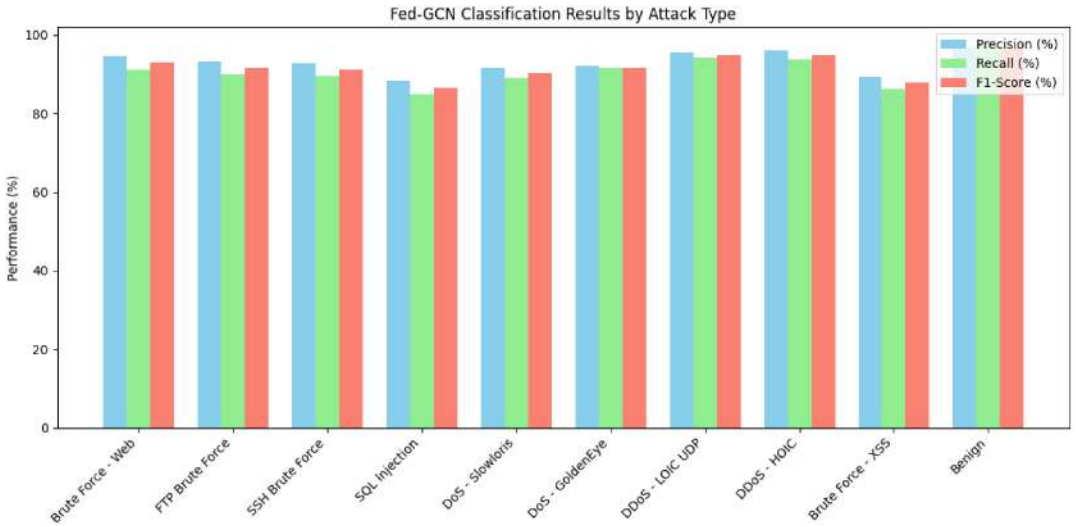

FIGURE 3 | Classification results per attack type.

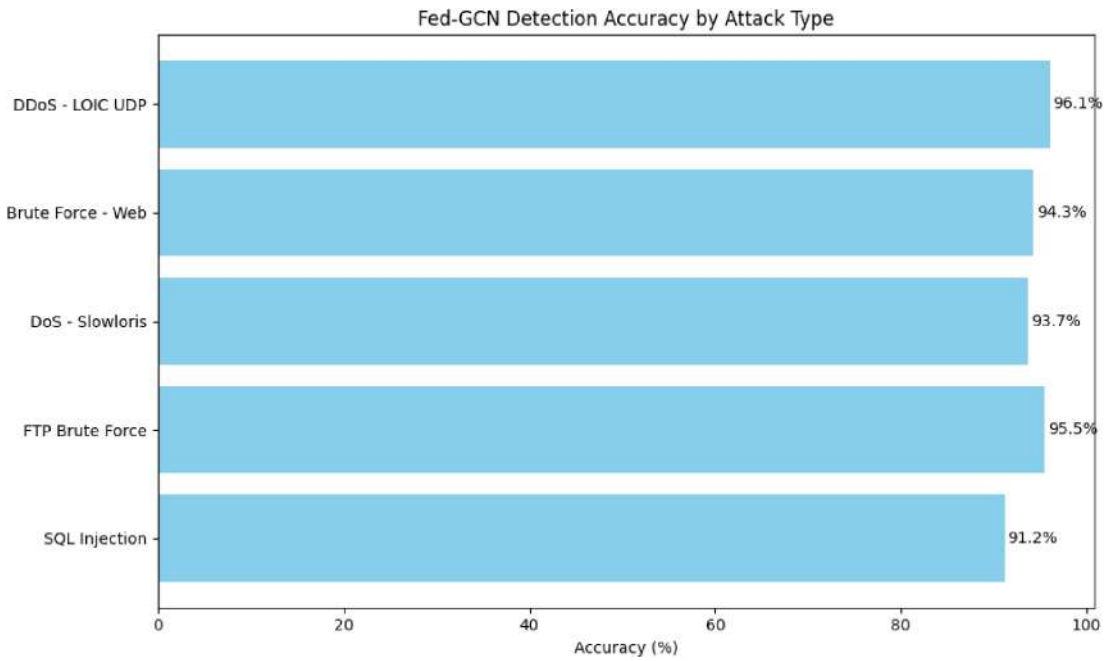

FIGURE 4 | Detection performance by attack type.

accuracy, precision, recall, F1-score, and false positive rate (FPR), which provide a clear picture of each model type's performance in detecting and classifying cyberattacks against IoT systems.

The Random Forest model achieves an accuracy of 90.4%, relatively moderate precision (88.9%) and recall (87.1%), and an F1-score of 88.0%. It exhibits a false-positive rate of 11.2%, characterized by a certain degree of real-time misclassification of benign traffic as threats. CNN is a little better, achieving 91.7% accuracy and a more balanced precision-recall curve (90.3% and 89.5%, respectively), but still a false-positive rate of 9.7%. This will create redundant alarms or system loading during actual deployment. To jump to graph-based architectures, the Traditional GCN outperforms Random Forest and CNN (achieving 93.2% accuracy, a 91.3% F1-score, and a lower false-positive rate of 8.5%) due to its ability to capture structural relationships in flow traffic data. Nevertheless, Fed-GCN easily outperforms all baselines by a wide margin. It achieves the best accuracy of 97.3, precision of 93.9, recall of 94.6, and F1-score of 94.2, and also minimizes the FPR to 6.7, compared to conventional centralized models, demonstrating a significant improvement. These findings establish that, in addition to improving detection accuracy, the Fed-GCN model can minimize misclassifications and that data privacy is not compromised by federated learning. Its better performance

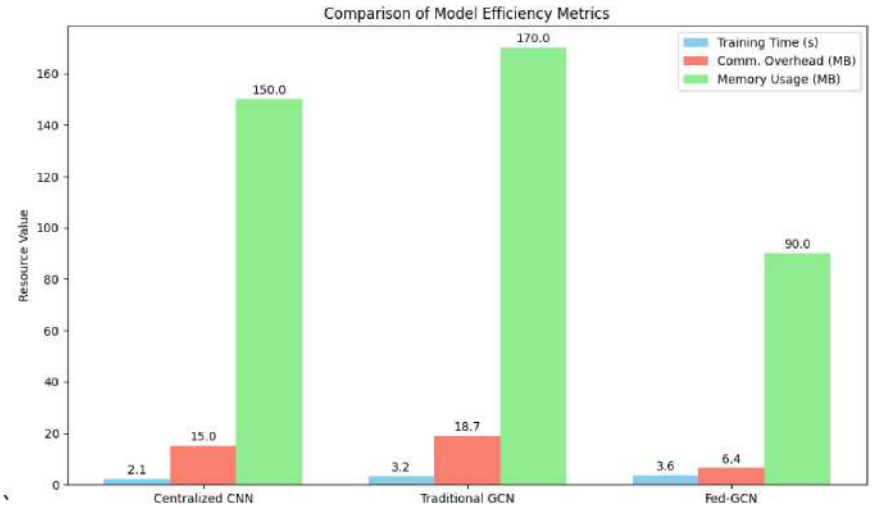

FIGURE 5 | Model efficiency and communication overhead.

TABLE 3 | Performance metrics – Fed-GCN vs. baseline methods.

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)	False positive rate (FPR) (%)
Random Forest	90.4	88.9	87.1	88.0	11.2
CNN	91.7	90.3	89.5	89.9	9.7
Traditional GCN	93.2	91.6	91.0	91.3	8.5
Proposed Fed-GCN	97.3	93.9	94.6	94.2	6.7

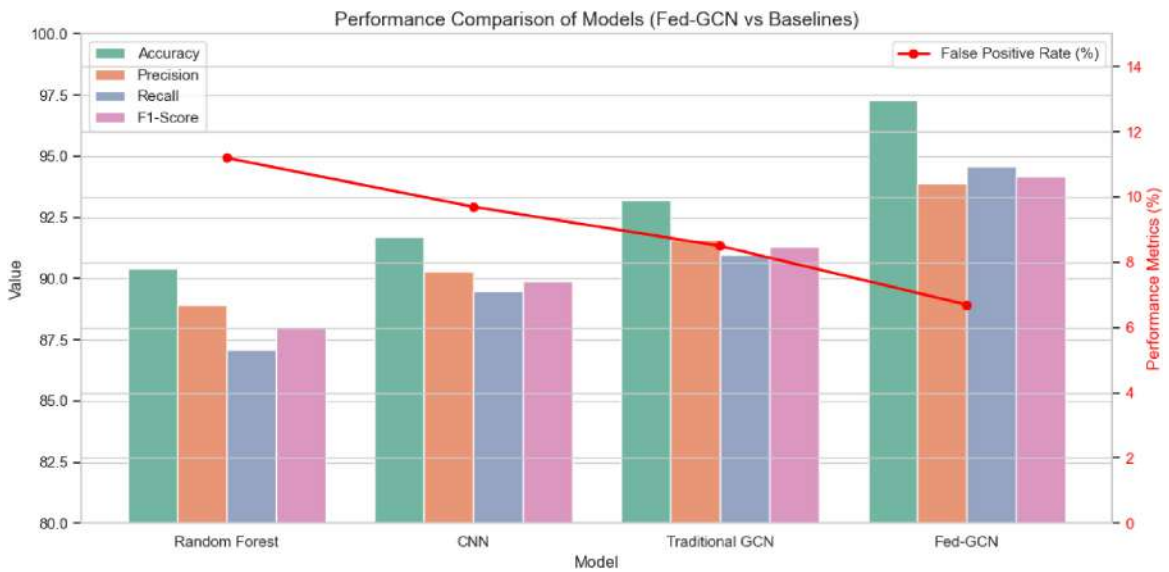

FIGURE 6 | Performance comparative assessment.

shows that the combination of graph neural networks and federated learning is effective for scalable, secure, and privacy-sensitive cyber threat detection in IoT networks. Moreover, the results of the performance comparative assessment are shown in **Figure 6**.

Discussion

The proposed Fed-GCN system is an innovation in IoT forensic intelligence, combining graph learning with federated privacy-preserving computation. A Fed-GCN representation of network flows through graphs can capture their complex interdependencies and thus detect distributed

attacks such as DDoS and brute-force attacks. This situational knowledge is fundamental to forensic investigation, as single anomalies can be meaningless unless they are put into context and/or compared with other traffic states. With the framework, federated learning is employed for on-device training, which implies that no transfer of sensitive raw data is necessary. The practice helps comply with data protection laws such as GDPR, enhances data sovereignty, and reduces the risk of data exposure during training. Although the controlled experiments demonstrated high accuracy and efficiency, there are significant challenges in deploying Fed-GCN in real-world IoT networks. Slow latency between the edge nodes and the central aggregator may delay timely model updates, particularly in environments with bandwidth constraints. The diversity of devices, such as processing power, memory, and operating systems, can also influence model convergence and consistency across nodes. Also, federated learning provides better privacy; however, clients may be compromised by poisoning or inference attacks. To ensure evidentiary certainty, the model uses blockchain-based logging to timestamp and securely, irreversibly store forensic data. However, this can introduce latency (one of the disadvantages of tamper-resistant evidence trails), particularly in high-throughput settings where forensic logging must scale as quickly as possible. Such trade-offs will also need to be considered to render Fed-GCN practical in the context of operational security architectures.

Conclusion

This paper proposes a new Fed-GCN architecture that enhances real-time forensics in the IoT setting. To address the existing critical need for decentralized, privacy-preserving, and context-aware cyber forensics, Fed-GCN utilizes graph structures to represent dependencies within the analyzed traffic flow and federated learning to keep sensitive data local. This method will also enable the organization to comply with global privacy laws such as GDPR and facilitate scalable, secure, and collaborative threat detection across various IoT nodes. The proposed model demonstrated stability and near-perfect accuracy across multiple attack types. Compared to traditional classifiers and centralized deep learning, the Fed-GCN detected a large number of causalities with much higher accuracy (97.3%). The Fed-GCN not only has higher detection accuracy than traditional classifiers but also achieves an overall F1-score of 94.2% and a drastically reduced FPR of 6.7%. Detection performance at the attack-specific level was good, and the model performed well at detecting DDoS-LOIC UDP (96.1%), FTP brute force (95.5%), and brute force-web (94.3%), among others. Also, forensic logging with Fed-GCN achieved 100% evidence integrity verification, 96.4% completeness, and close-to-real-time event logging with an average of 128 ms per event, indicating that Fed-GCN is ready for deployment in a real-time system. The communication overhead was small compared to the centralized models, at just 6.4 MB per round, and the model was memory-efficient, with a 90 MB footprint, making it appropriate for resource-constrained edge devices. In the future, the system will be integrated with edge-based AI accelerators to lower latency and power consumption, adaptive learning to address evolving and even zero-day threats, and policy-based compliance to meet regulatory requirements in industries such as healthcare and critical infrastructure. The supplemental components will also improve the forensic soundness, legal admissibility, and practical feasibility of the Fed-GCN framework in practical IoT environments. The study's terminal result is a scalable foundation for intelligent, compliant, and proactive cybersecurity in the ever-more interdependent digital world.

Funding

The research received no external funding, and this has been explicitly stated in the paper.

Conflict of Interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

1. Gillis AS, Yasar K. *What is IoT (Internet of Things)? | Definition from TechTarget*. IoT (2025). Available online at: https://www.techtarget.com/iotagenda/definition/Internet-of-Things-IoT (Accessed July 31, 2025).
2. Al-Sarawi S, Anbar M, Abdullah R, Al Hawari AB. Internet of things market analysis forecasts, 2020–2030. In: *2020 Fourth World Conference on Smart Trends in Systems, Security and Sustainability (WorldS4)*. London, United Kingdom: IEEE (2020). p. 449–53. doi: 10.1109/WorldS450073.2020.9210375
3. Sasi T, Lashkari AH, Lu R, Xiong P, Iqbal S. A comprehensive survey on IoT attacks: taxonomy, detection mechanisms and challenges. *J Inf Intell*. (2024) 2(6):455–513. doi: 10.1016/j.jiixd.2023.12.001
4. Gelgi M, Guan Y, Arunachala S, Samba Siva Rao M, Dragoni N. Systematic literature review of IoT botnet DDOS attacks and evaluation of detection techniques. *Sensors*. (2024) 24(11):11. doi: 10.3390/s24113571
5. Qureshi SU, He J, Tunio S, Zhu N, Nazir A, Wajahat A, et al. Systematic review of deep learning solutions for malware detection and forensic analysis in IoT. *J King Saud Univ Comput Inf Sci*. (2024) 36(8):102164. doi: 10.1016/j.jksuci.2024.102164
6. Papadopoulos C, Kollias K-F, Fragulis GF. Recent advancements in federated learning: state of the art, fundamentals, principles, IoT applications and future trends. *Future Internet*. (2024) 16(11):11. doi: 10.3390/fi16110415
7. Elmisery AM, Sertovic M. Enabling collaborative forensic by design for the internet of vehicles. *Information*. (2025) 16(5):5. doi: 10.3390/info16050354

8. Igonor OS, Amin MB, Garg S. The application of blockchain technology in the field of digital forensics: a literature review. *Blockchains*. (2025) 3(1):1. doi: 10.3390/blockchains3010005
9. Makura S, Venter HS, Kebande VR, Karie NM, Ikuesan RA, Alawadi S. Digital forensic readiness in operational cloud leveraging ISO / IEC 27043 guidelines on security monitoring. *Secur Priv*. (2024) 4(3):e149. doi: 10.1002/spy2.149
10. Sadineni L, Pilli ES, Battula RB. Ready-IoT: a novel forensic readiness model for internet of things. In: *2021 IEEE 7th World Forum on Internet of Things (WF-IoT)*. New Orleans, LA, USA: IEEE (2021). p. 89–94. doi: 10.1109/WF-IoT51360.2021.9595902
11. Kebande VR, Menza NK, Venter HS. Functional requirements for adding digital forensic readiness as a security component in IoT environments. *Int J Adv Sci Eng Inf Technol.* (2018) 8(2): 342.
12. Imperva. *Cybersecurity Threats | Types & Sources*. Imperva (2025). Available online at: https://www.imperva.com/learn/application-security/cyber-security-threats/ (Accessed July 31, 2025).
13. Ghelani D. Cyber security, cyber threats, implications and future perspectives: a review. *Authorea Preprints*. (2022). Available online at: https://www.authorea.com/doi/full/10.22541/au.166385207.73483369 (Accessed July 31, 2025).