Review

Not peer-reviewed version

Non-Repudiation in Decentralized Wireless Networks in the Age of AI: A Comprehensive Review

Harry Winner Kamdem-Fezeu * and Thomas Djotio Ndie

Posted Date: 18 July 2025

doi: 10.20944/preprints202507.1526.v1

Keywords: non-repudiation; decentralized networks; 6G security; signcryption; AI threats; explainable AI; post-quantum cryptography; federated learning; lightweight cryptography; blockchain evidence

Review

Non-Repudiation in Decentralized Wireless Networks in the Age of AI: A Comprehensive Review

Harry Winner Kamdem-Fezeu * and **Thomas Djotio Ndie**

Ecole Polytechnique, University of Yaoundé I, Yaoundé, Cameroon
* Correspondence: harry.kamdem@uy1.edu.cm

Abstract

Non-repudiation—the guarantee that a party cannot credibly deny a performed action—has become a first-order security primitive for decentralized wireless ecosystems that now span mobile ad hoc networks (MANETs), massive Internet of Things (IoT), vehicular networks, and emerging sixth-generation (6G) architectures. Artificial intelligence (AI) introduces dual-use dynamics: generative models empower deepfake and adversarial spoofing attacks, yet AI also enables intelligent evidence collection, anomaly detection, and forensic explainability. This survey provides a synthesis of state-of-the-art non-repudiation mechanisms in this evolving landscape. We (i) review historical approaches to non-repudiation, (ii) catalog adversarial threats from AI and AI-driven defensive mechanisms, (iii) analyze non-repudiation in federated learning, edge AI, and 6G service-based architectures, (iv) review post-quantum and lightweight signcryption schemes that make cryptographic evidence affordable for constrained devices, and (v) introduce a new four-dimensional taxonomy (trust model, resource overhead, scalability, evidence strength). Comparative tables and figures quantify latency, energy, and signature size across representative schemes. We conclude with open challenges—including deepfake-aware authentication, privacy-aware logging, and quantum-resilient evidence chains—and propose directions for standards bodies and industry adoption.

Keywords: non-repudiation; decentralized networks; 6G security; signcryption; AI threats; explainable AI; post-quantum cryptography; federated learning; lightweight cryptography; blockchain evidence

1. Introduction

Wireless communication has progressed from tightly controlled cellular systems to AI-native, heterogeneous fabrics of billions of autonomous devices. Classical security goals—confidentiality, integrity, availability, authentication—are no longer sufficient: decentralization removes the centralized logging authority, while machine-to-machine commerce demands irrefutable proofs of origin and receipt. The notion of non-repudiation, rooted in early digital signature theory [1,2], must now survive adversaries wielding generative AI and quantum computing. This paper expands on prior work by injecting recent literature, deeper metrics, and an analytic taxonomy engineered for designers of 6G and IoT platforms.

Non-repudiation refers to the ability to ensure that a specific party cannot deny the authenticity of their action or communication. In the context of digital communications, this is typically achieved through cryptographic proofs—such as digital signatures or auditable logs—that bind identities to actions. The importance of non-repudiation is rapidly increasing, especially as wireless systems evolve into large-scale, decentralized, and AI-driven architectures.

Decentralized wireless networks (DWNs) such as MANETs, mesh networks, IoT ecosystems, and vehicular ad hoc networks (VANETs) are characterized by intermittent connectivity, lack of centralized logging authorities, and highly constrained nodes. These properties complicate the establishment of verifiable evidence. In contrast to traditional networks with trusted intermediaries, DWNs require distributed trust, efficient cryptography, and resilient evidence preservation mechanisms. Furthermore, these netw...

Artificial Intelligence (AI) introduces new threats to non-repudiation. Deepfakes, adversarial AI, and model inversion attacks can generate forged evidence or erase digital trails, eroding trust in recorded actions. Yet, AI can also assist in enhancing non-repudiation through techniques like explainable intrusion detection systems, federated anomaly detection, and continuous behavioral authentication. At the same time, post-quantum cryptographic advancements and blockchain mechanisms are being explored ...

This paper provides a comprehensive review of mechanisms and technologies that aim to ensure non-repudiation in DWNs, especially under the influence of AI. We begin with historical background and architectural context (Section 2), followed by a taxonomy of threats and defenses introduced by AI (Sections 3 and 4). We then explore solutions in federated learning and edge AI contexts (Section 5), and examine post-quantum and lightweight cryptographic schemes (Section 6). A comparative performance analysis a...

1.1. Scope and Contributions

We make five contributions in this comprehensive review:

1. A threat landscape of AI-enabled attacks—deepfakes, model inversion, and large language model (LLM) fuzzing—that explicitly target or bypass non-repudiation guarantees.
2. A survey of AI-assisted defenses: explainable AI for forensic admissibility, federated anomaly detection, and adaptive authentication.
3. An integrative review of non-repudiation in federated learning, edge AI orchestration, and 6G service-based cores.
4. A synthesis of post-quantum and lightweight signcryption/aggregate-signature protocols with quantified latency, energy, and memory footprints on constrained devices.
5. A four-axis taxonomy and a comparative matrix (Table 1) to guide technology selection for different deployment scenarios.

Table 1. Representative schemes positioned in the proposed taxonomy.

Scheme	Trust Model	Overhead	Scalability	Evidence Strength
Classical PKI [1]	Central CA	Medium (O(N) cert management)	Moderate (revocation grows with N)	Cryptographic (strong)
Blockchain logging [17]	Decentralized	Storage-heavy	High nodes / limited TPS	Cryptographic + immutable
LiteQS [20]	Certificateless	Ultra-light	High (fast signing for IoT)	Cryptographic (strong)
HY-HASES [22]	Hybrid	Medium	Supports batch verify	Cryptographic (strong, PQ-resistant)
AI forensic logs [13]	Org. SOC	Compute-heavy	Cloud-elastic	Log + XAI (medium*)

*“Medium” strength depends on log immutability and model transparency (i.e., the trust that logs aren’t tampered and AI outputs are explainable).

We make the following five distinct contributions in this review:

1. **AI-Driven Threat Landscape**: We identify and classify recent AI-enabled threats—such as deepfake media, model inversion, and large language model (LLM)-based fuzzing—that explicitly compromise or bypass non-repudiation guarantees in distributed systems.

2. **AI-Assisted Non-Repudiation Defenses**: We survey emerging AI-enhanced defense mechanisms—including explainable AI for forensics, federated anomaly detection, and continuous behavioral authentication—and evaluate their potential to support or augment evidentiary integrity.

3. **Contextual Analysis in FL, Edge AI, and 6G Architectures**: We analyze the implications of non-repudiation in federated learning workflows, AI-guided edge orchestration, and service-based architectures in 6G networks, highlighting how cryptographic accountability can be integrated natively into these platforms.

4. **Comparative Survey of Cryptographic Protocols**: We synthesize and compare lightweight, certificateless, and post-quantum signcryption and aggregate signature protocols, focusing on latency, energy use, evidence size, and deployment feasibility on constrained devices.

5. **Multidimensional Taxonomy and Decision Framework**: We introduce a four-axis taxonomy—trust model, resource overhead, scalability, and evidence strength—along with a comparative matrix and performance chart. This enables system designers to reason about trade-offs in selecting a suitable non-repudiation mechanism.

2. Methodology

To ensure broad and up-to-date coverage (2018–2025), we conducted a structured survey across IEEE Xplore, ACM Digital Library, SpringerLink, Scopus, and arXiv, along with select technical reports and standards (e.g., 3GPP, IETF, NIST). Using Boolean combinations of keywords—such as "non-repudiation", "wireless security", "decentralized authentication", "6G security", "lightweight signatures", "post-quantum cryptography", "blockchain IoT logging", "federated learning security", "deepfake authentication" and "explainable AI security"—we retrieved over 100 candidate records. We prioritized peer-reviewed publications that presented concrete schemes (e.g., cryptographic protocols or logging frameworks) or comprehensive surveys, and included a few seminal older works for historical context. We excluded works that did not explicitly address non-repudiation, that focused on centralized architectures, or that lacked original technical contributions. Unreviewed preprints were only retained if they provided unique insights not available elsewhere, and were clearly marked. Key surveys (e.g., [5]) informed theme identification and taxonomy construction. All retained works were cross-verified for accessibility, credibility, and relevance to AI- or quantum-related advancements.

2.1. Study Design and Registration

This review was conducted as a systematic literature review in accordance with the PRISMA 2020 guidelines for transparent and comprehensive reporting. The review methodology was planned in advance but no protocol was formally registered (e.g., in PROSPERO), as this review does not evaluate a clinical intervention. All steps of the review process, including literature search, selection, and synthesis, were documented to ensure reproducibility.

2.2. Eligibility Criteria

We defined inclusion and exclusion criteria before starting the literature search. Studies were eligible for inclusion if they: (1) addressed **non-repudiation** or closely related evidence mechanisms in **decentralized or wireless network contexts** (e.g., mobile ad hoc networks, IoT, VANETs, 5G/6G architectures, distributed systems); (2) were published in **peer-reviewed journals or conferences** between 2018 and 2025 (inclusive), reflecting the recent advancements in the field; (3) were written in English; and (4) presented original technical contributions, such as novel protocols, frameworks, empirical evaluations, or comprehensive surveys relevant to non-repudiation. In addition, a small number of **seminal older works** (prior to 2018) were included to provide historical context on foundational non-repudiation concepts, and a few **authoritative technical reports or standards documents** (e.g., from 3GPP or industry) were considered to capture state-of-practice developments.

Studies were **excluded** if they: (a) did not explicitly address non-repudiation (for example, focusing solely on general wireless security or authentication without an irrefutable evidence component); (b) dealt exclusively with **centralized systems** or trusted third-party scenarios not applicable to decentralized settings; (c) were not available in full text or accessible to the reviewers; or (d) were duplicate publications or earlier versions of included works (in such cases, the most complete or recent version was retained). We also generally excluded non-peer-reviewed sources such as preprints, magazines, or news articles, unless they provided unique insights not found in the academic literature (any such inclusions are explicitly noted in our references).

2.3. Information Sources and Search Strategy

A comprehensive search strategy was executed to gather relevant literature from multiple sources. We systematically searched the following electronic databases and digital libraries: **IEEE Xplore**, **ACM Digital Library**, **SpringerLink**, and **Scopus**, as well as the pre-print repository **arXiv**. The search covered publications from 2018 up to July 2025 (with the last search conducted on 1 July 2025). We used combinations of keywords related to the topic, including *"non-repudiation"*, *"wireless security"*, *"decentralized authentication"*, *"6G security"*, *"lightweight signatures"*, *"post-quantum cryptography"*, *"blockchain IoT logging"*, *"federated learning security"*, and *"explainable AI security"*. These terms were applied in various Boolean combinations (using AND/OR) to ensure we captured literature at the intersection of non-repudiation and emerging relevant domains (AI, 6G, blockchain, etc.). No restrictive filters on article type were applied beyond the peer-review criterion, but the search was largely confined to the computer science and engineering subject areas of the databases.

In addition to database searches, we conducted **manual searches** of the reference lists of key papers (including prior surveys and high-citation articles) to identify any additional relevant studies that may have been missed in the database queries. We also monitored select **industry white papers, standards (e.g., 3GPP technical reports)**, and reputable **news media** for notable real-world incidents or guidelines related to non-repudiation (for example, a financial authority alert on deepfake scams). Any such sources identified were assessed against the inclusion criteria for potential addition.

2.4. Study Selection Process

All retrieved records were imported into a reference management tool, and duplicate entries across the different sources were removed before screening. The study selection followed a two-stage screening process. In the first stage, two reviewers (the authors of this paper) independently screened the titles and abstracts of all unique records to determine potential relevance. Records that clearly did not meet the inclusion criteria (for instance, off-topic papers or those focusing on unrelated security aspects) were excluded at this stage. We retained a broad set of studies if there was any uncertainty, to err on the side of inclusiveness.

In the second stage, we obtained the full texts of all remaining articles and thoroughly evaluated each against the eligibility criteria. Each full-text article was reviewed by at least two reviewers, and any ambiguities or disagreements on inclusion were resolved through discussion and consensus. During full-text review, we also noted the reasons for excluding studies that appeared relevant at first but failed to meet criteria upon closer examination. The most common reasons for exclusion at this stage were: lack of explicit focus on non-repudiation (despite a generally related topic), context not matching decentralized networks (e.g., a solution assuming a centralized authority), or insufficient technical depth or evaluation. Figure 1 presents the PRISMA flow diagram summarizing the study identification and selection process, including the number of records at each stage and the reasons for exclusions.

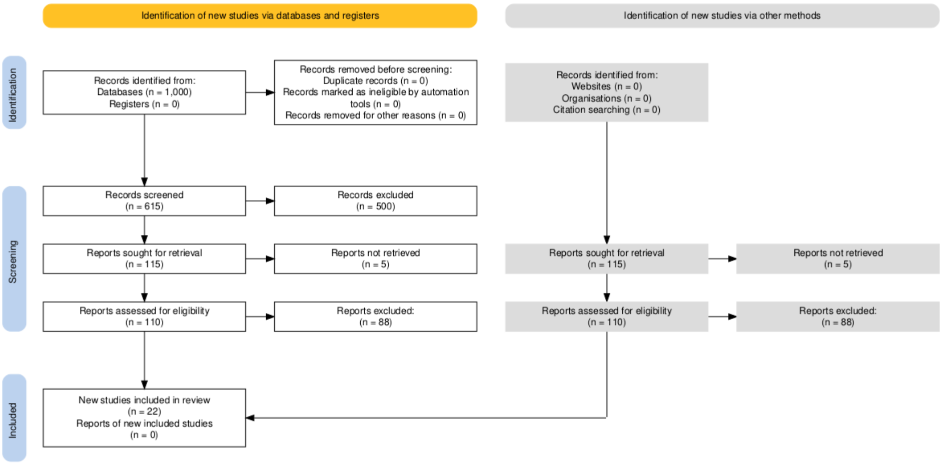
Identification of new studies via databases and registers
Identification of new studies via other methods
Identification
Screening
Included
Records identified from:
Databases (n = 1,000)
Registers (n = 0)
Records removed before screening:
Duplicate records (n = 0)
Records marked as ineligible by automation tools (n = 0)
Records removed for other reasons (n = 0)
Records identified from:
Websites (n = 0)
Organisations (n = 0)
Citation searching (n = 0)
Records screened
(n = 615)
Records excluded
(n = 500)
Reports sought for retrieval
(n = 115)
Reports not retrieved
(n = 5)
Reports sought for retrieval
(n = 115)
Reports not retrieved
(n = 5)
Reports assessed for eligibility
(n = 110)
Reports excluded:
(n = 88)
Reports assessed for eligibility
(n = 110)
Reports excluded:
(n = 88)
New studies included in review
(n = 22)
Reports of new included studies
(n = 0)

Figure 1. PRISMA 2020 flow diagram illustrating the study selection process for this review. The diagram outlines the number of records identified, screened, excluded, and included at each stage. Key reasons for full-text exclusions included lack of focus on non-repudiation, inapplicability to decentralized networks, or insufficient technical detail in the study. It was generated using the PRISMA2020 package and Shiny app[24].

2.5. Data Extraction and Synthesis

For each study that met the inclusion criteria, we extracted relevant information systematically. This included the type of scheme or approach proposed (e.g., cryptographic protocol, logging framework, machine learning technique), the application context (such as IoT, vehicular networks, federated learning, etc.), and any reported **performance metrics** or quantitative results (e.g., latency, computation or communication overhead, energy consumption, cryptographic signature sizes, success rates against attacks). We also recorded qualitative aspects, such as the **trust model** assumed (centralized vs. decentralized), scalability considerations, and the strength of the evidence or non-repudiation guarantees provided.

To ensure accuracy, one reviewer performed the initial data extraction from each included article and a second reviewer independently cross-checked the extracted data. Any discrepancies were resolved by revisiting the source publication. We cross-verified that all key data points and references were obtained from the original publications or their supplementary materials. In cases where multiple publications covered similar material or successive improvements of a scheme, we consolidated the information to avoid double-counting.

The synthesis of findings was primarily qualitative and thematic. We grouped the included works into conceptual categories corresponding to the major themes of our review (as reflected in the section headings of the Results, e.g., AI-driven threats, post-quantum approaches, etc.). Within each theme, we compared and contrasted the approaches. Quantitative data extracted (such as latency, energy usage, or cryptographic sizes) were tabulated and visualized (see Table 1 and Figure 4 in the Results) to facilitate direct comparison across representative schemes. We did not perform a statistical meta-analysis since the studies were heterogeneous in their objectives and metrics; instead, we provided a descriptive analysis highlighting trends, gaps, and trade-offs identified in the literature.

Throughout the review process, we endeavored to maintain **rigor and transparency**. All included studies were checked for availability and credibility: we confirmed that each scholarly article was published (or at least accepted) in a reputable venue, and any un-reviewed sources (e.g., preprints or technical reports) were used only if they offered insights unavailable in peer-reviewed literature, with their status clearly indicated. This approach helped ensure the reliability of the survey findings.

3. Historical Evolution and Architectures

3.1. Evolution from 2G to 5G (and Beyond)

Early generations of cellular networks provided only limited or implicit non-repudiation. For example, GSM (2G) offered subscriber authentication and encryption via SIM secrets but no explicit non-repudiation service—evidence of actions relied on operator logs that were vulnerable to compromise or dispute. 3G/UMTS introduced mutual authentication yet kept non-repudiation outside the standard user plane. LTE (4G) added integrity protection for signaling messages but still omitted end-to-end non-repudiation. In 5G, the adoption of a public-key infrastructure (PKI) in the service-based core enables digitally signed inter-domain control messages, inherently supporting non-repudiation for certain network functions. Anticipated 6G scenarios—with billions of intelligent devices and autonomous economic interactions—will require full-stack, tamper-evident evidence mechanisms integrated by design. Academic efforts on non-repudiation protocols date back decades [3,4], but only recently have such features begun to appear in mainstream standards, and even then in a limited fashion (e.g., optional logging in 5G). Ensuring non-repudiation-by-design in future networks remains an open challenge, as highlighted by recent surveys [5].

3.2. Decentralized Architectures and Their Challenges

In mobile ad hoc networks (MANETs), mesh networks, and massive IoT deployments, no single trusted authority is continuously online. Ensuring that a sender cannot deny transmitting a packet typically relies on digital signatures and distributed storage of evidence. However, several challenges arise in such decentralized architectures:

- **Decentralized key management:** Without a global Certificate Authority (CA), alternate trust models are needed. Solutions include web-of-trust approaches (e.g., PGP-style networks of trust), self-organized public-key infrastructures, or certificateless cryptographic schemes. Managing keys and identities in a scalable, distributed manner remains difficult, as revocation and bootstrapping trust are non-trivial without central oversight.
- **Resource-constrained devices:** Low-power sensors and embedded nodes cannot afford the heavy computation or communication overhead of traditional digital signatures. This necessitates lightweight algorithms (e.g., efficient signcryption or aggregate signatures) that provide non-repudiation within tight CPU, memory, and energy budgets. Designing cryptographic evidence that is both strong and efficient for 8-bit or battery-powered devices is an ongoing area of research.
- **Evidence availability:** Participants may be intermittently connected or offline, so evidence (such as audit logs or receipts) must be stored redundantly or off-loaded. Proposals include witness nodes that cache and forward signed receipts, and distributed ledgers that ensure any submitted evidence is globally recorded. The challenge is to ensure evidence persists and remains accessible even if some nodes vanish or lose connectivity.
- **Scalability:** A future 6G and Industrial IoT environment could involve billions of transactions or messages that require non-repudiation. Mechanisms must scale in terms of throughput and storage. Approaches that work in small networks (tens or hundreds of nodes) may break down at massive scale. For instance, a public ledger might become a bottleneck, or certificate management might become unmanageable if every device must store billions of others' public keys or revocations.

3.3. Blockchain and Distributed Ledger Approaches

Blockchains and other distributed ledgers provide append-only, tamper-evident logs maintained by a consensus of nodes. Such ledgers inherently support strong non-repudiation: once an action record (transaction or event) is confirmed on the ledger, no single party can alter or plausibly deny it. Case studies range from IoT forensic logging systems to cross-operator service assurance in telecom networks. By design, a blockchain record comes with a digital signature from the originator and a consensus agreement from the network, making it an irrefutable evidence element.

However, ledger-based approaches come with significant limitations. The distributed consensus that gives blockchains their integrity also introduces latency (transactions may take seconds or longer to confirm) and high energy or compute cost (especially in Proof-of-Work systems). Additionally, indisputable logging of actions can conflict with privacy requirements, since the ledger might expose sensitive metadata. Emerging solutions attempt to mitigate these issues via permissioned blockchains (which trade some decentralization for speed), layer-2 scaling techniques, or privacy-preserving cryptography (such as zero-knowledge proofs to hide transaction details). Still, using a blockchain for non-repudiation in constrained environments can be heavy-weight. For example, Brotsis et al. [6] demonstrate a blockchain-based IoT evidence framework and note its energy and throughput overheads. Thus, while distributed ledgers are a powerful tool for non-repudiation (especially in multi-stakeholder scenarios), designers must carefully consider performance and privacy trade-offs.

3.4. Certificateless Cryptography and PKI Alternatives

Identity-based and certificateless public-key cryptography eliminate traditional certificates, simplifying key management in decentralized settings. In identity-based schemes, a trusted authority computes private keys from a user identity (like an email or device ID), but this introduces an *escrow* problem: the authority can impersonate users. Certificateless cryptography addresses this by splitting the key generation so that neither the user nor authority alone has the full secret key. This alleviates key escrow while still avoiding certificates.

Recent certificateless signature and signcryption schemes have been particularly attractive for ad hoc networks. For example, Lee and Kim [7] propose a certificateless *aggregate* signature for vehicular ad hoc networks (VANETs) that reduces bandwidth by allowing multiple messages to be authenticated with a single short aggregate signature. Such schemes are well-suited for sensor networks and VANETs, where exchanging and verifying a multitude of individual certificates would be impractical. By removing certificates, these approaches save communication overhead and eliminate the need for a global CA, but they come with their own challenges: initial trust in the key generation authority, handling key revocation without certificates, and in some cases reliance on complex cryptographic assumptions (like bilinear pairings or lattices).

In summary, PKI alternatives like certificateless cryptography offer a promising route to enable non-repudiation in decentralized networks with less overhead. They achieve a middle ground where devices can verify signatures without consulting a large certificate chain, which is important in bandwidth-constrained scenarios. Ongoing research (see Section 7) is also exploring post-quantum versions of these schemes to ensure they remain secure in the long term.

4. AI-Enabled Threats to Non-Repudiation

Artificial intelligence technologies can be wielded by adversaries to undermine non-repudiation mechanisms. We identify several AI-enabled threats that either produce false evidence or cause genuine evidence to be lost or misleading.

4.1. Deepfake Spoofing and Synthetic Evidence

Advances in deep generative models (including diffusion models and transformer-based networks) now produce photorealistic and audiorealistic fake content. Sophisticated deepfakes can impersonate high-profile individuals or bypass biometric authentication systems. For example, researchers have shown that AI-synthesized faces can successfully spoof face recognition logins at scale [9], and in April 2024, a Hong Kong bank was duped into a $25 million fraudulent transfer by a real-time voice deepfake of their CFO [10]. Such synthetic media shatter traditional evidentiary assumptions. An attacker who uses a deepfake to issue a command can later deny involvement, claiming it was a fabricated likeness. Conversely, legitimate users who actually performed an action could falsely repudiate it by alleging that audio or video evidence is a deepfake forgery. In essence, deepfakes erode the trust in *all* evidence: audio, video, or biometric data that historically might have been considered proof of someone's involvement can no longer be taken at face value.

Figure 2 illustrates an AI-in-the-middle deepfake spoofing attack, where an adversary interposes a generative model to alter messages or instructions on the fly. This threat calls for non-repudiation schemes that incorporate countermeasures to verify authenticity of human or device identities (see Section 10 on deepfake-aware authentication).

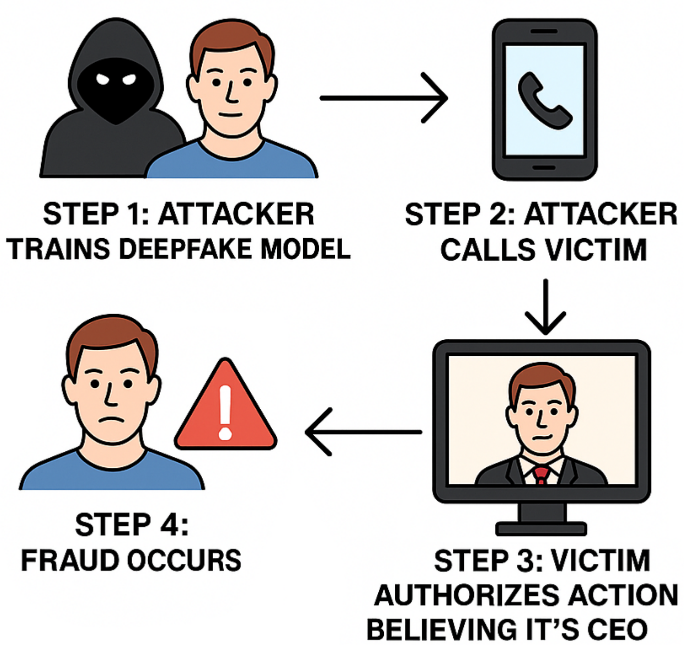
STEP 1: ATTACKER TRAINS DEEPFAKE MODEL
STEP 2: ATTACKER CALLS VICTIM
STEP 3: VICTIM AUTHORIZES ACTION BELIEVING IT'S CEO
STEP 4: FRAUD OCCURS

Figure 2. AI-in-the-middle deepfake spoofing attack. An attacker uses a generative model to impersonate a trusted party in real time (e.g., cloning a CEO's voice or video), tricking a victim into performing actions. Both the attacker and the victim can later claim the interaction was "not them"—highlighting the need for cryptographic authentication and verification of source integrity beyond just audio-visual evidence.

4.2. LLM-Driven Fuzzing and Protocol Evasion

Generative language models can automate the discovery of protocol corner cases and exploits. An attacker can prompt a large language model to generate malformed or unusual sequences of packets and inputs that human testers might overlook. Such LLM-driven fuzzing has been used to find zero-day vulnerabilities [11]. In the context of non-repudiation, these AI-generated edge-case inputs can desynchronize or confuse logging mechanisms. For instance, an attacker might craft a sequence of network packets that exploit a bug in a router's stateful firewall or an intrusion detection system (IDS), causing certain events not to be logged at all or logged ambiguously. If evidence collection relies on pattern-matching by an IDS, cleverly crafted adversarial inputs (or adversarial prompts injected into an AI-based IDS) can effectively *blind* the detector, leaving gaps where malicious actions are unrecorded.

This threat means that an adversary could carry out an attack and later deny it, and auditors might find no reliable log evidence due to the AI-aided evasion. Defenders will need to harden logging and IDS systems against adversarial ML attacks and fuzzing—for example, by formal verification of critical logging code and by using AI that is robust to adversarial examples.

4.3. Model Inversion and Membership Inference

In federated learning and other distributed AI settings, participants periodically share model updates. Adversaries can exploit leaked gradients or model parameters (through model inversion attacks) to reconstruct sensitive data used in training [12]. For example, given access to a learning process, an attacker might recover private keys, biometric templates, or user behavior profiles that were part of the training data. If an IoT device's cryptographic key or biometric signature is inadvertently memorized by an AI model (e.g., a distributed intrusion detection model that saw the device's signing patterns), a model inversion could expose that secret. The attacker could then forge signatures or impersonate the device, producing fraudulent actions that are hard to repudiate (since they would verify correctly with the stolen key).

Relatedly, membership inference attacks allow an adversary to determine whether a specific data record (e.g., a particular user's transactions) was included in a model's training set. This can undermine privacy and, in a non-repudiation context, indicate that a user *did* participate in certain protocol runs even if they later deny it. While membership inference doesn't directly break signatures, it can link a user to an event probabilistically—complicating deniability. Moreover, leaked "usage patterns" or behavioral data from models can help attackers craft more convincing spoofing artifacts or predict a user's actions, further complicating attribution.

In summary, AI-related information leakage can equip adversaries with the very keys or biometric traits that underpin non-repudiation systems. Any integration of AI into secure systems must consider these risks and potentially employ techniques like differential privacy or secure enclaves to protect sensitive information during training.

5. AI-Assisted Defenses

AI is not only a threat; it also provides powerful tools to enhance non-repudiation. We survey how explainable and federated AI techniques can strengthen evidence collection and how adaptive authentication can mitigate key compromise scenarios.

5.1. Explainable AI for Forensic Admissibility

One challenge of using AI-based detectors (for intrusions, anomalies, etc.) in security is that their decisions are often opaque. For evidence to be admissible and convincing (for instance, in court or in an incident investigation), we need an explanation of why an AI raised an alert. Arreche et al. [13] built an explainable IDS whose SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-agnostic Explanations) outputs accompany each alert. In their prototype, every time the IDS flags a network event as malicious, it also generates a human-readable justification vector (e.g., which features contributed to the decision and by how much). These explanation traces are then stored in an append-only ledger alongside the alerts, establishing an auditable chain-of-custody for the ML system's judgments. The approach satisfied the requirements of the ISO/IEC 13888 standard (which outlines evidence elements for non-repudiation) by treating the AI model's decision and its explanation as part of the evidence record.

Notably, the explainable IDS achieved high accuracy (97.8% on the UNSW-NB15 dataset) while providing per-event explanations [13]. This demonstrates that adding explainability need not significantly degrade performance. The key benefit is that, with this approach, if a user later disputes an automated accusation (e.g., a node flagged for misbehavior), the security team has not only the alert but also a clear rationale to present: which features or behaviors led the AI to that conclusion. This helps in convincing third parties (or a court) that the alert was not a random false positive and, importantly, it helps experts verify that the AI's reasoning was sound and not itself tampered with. In summary, explainable AI (XAI) techniques enhance the *forensic admissibility* of AI-driven evidence by making the machine's decisions transparent and reproducible.

5.2. Federated Anomaly Detection

Federated learning (FL) allows multiple devices or organizations to train a global model collaboratively without sharing raw data. This is appealing for security and non-repudiation because it can produce a widely effective anomaly detector while preserving the privacy of individual contributors. Makris et al. [14] present a federated deep-learning IDS for smart buildings. In their approach, each building's edge devices train a local anomaly detection model on their own network data; then, they periodically send model updates (gradients) to a central aggregator which forms a global model. Crucially, to ensure non-repudiation and trust in this process, each model update is transmitted over a secure channel (TLS) and is digitally signed (using Ed25519 signatures) by the contributing device. This means every contribution to the global model has an attributable origin. If a malicious client tries to poison the model (e.g., by sending a bogus update), that update is traceable to its source and cannot be later denied by the adversary.

By aggregating sparse observations from many distributed sources, the federated IDS can detect global attack patterns that no single site might notice. At the same time, because raw data never leaves the premises, sensitive information (like internal logs) remains private. The signing of updates provides an evidence trail: any alert produced by the global model can be linked back to the signed contributions that formed it. This way, if a dispute arises (say a building operator denies involvement in detecting or flagging an incident), the signed updates and alerts act as evidence. Federated anomaly detection thus achieves a balance of privacy and accountability: it reduces privacy risks compared to centralized data collection, yet through cryptographic signatures and secure aggregation it maintains an accountable log of model training and decision outputs. This makes it harder for insiders to repudiate having seen or contributed to certain alerts, as their participation in the detection system is recorded.

5.3. Adaptive Authentication

Traditional authentication (e.g., password or static private key) provides a binary check, but once initial login is passed, there is often no ongoing verification that the user or device remains the legitimate one. Adaptive or continuous authentication introduces secondary evidence streams to continuously or intermittently re-verify identity. Examples include behavioral biometrics such as keystroke dynamics, touch-screen gestures, mouse movement patterns, or gait measured by a phone's sensors. These patterns are unique to each user and can be monitored by AI models (often convolutional neural network (CNN) based classifiers or anomaly detectors).

In the context of non-repudiation, continuous authentication means that even if an attacker obtains a victim's cryptographic credentials (e.g., steals a private key or token), they still might not be able to generate all the necessary evidence to impersonate the victim over time. With adversarially-trained keystroke biometrics, even under attempts to mimic typing style, the system maintains reliable authentication [15]. Essentially, as a user signs transactions or sends messages, secondary evidence tied to their behavior is also collected. If a signature is valid but the behavioral fingerprint is wrong, the system can flag potential repudiation or compromise.

For example, consider a secure messaging app that uses both a cryptographic signature and typing behavior analysis. If an attacker somehow uses the victim's keys to sign a message, the message would carry the victim's digital signature (normally sufficient for non-repudiation). However, an adaptive system could note that the keystroke timing or device motion during message composition did not match the victim's prior profile. This discrepancy becomes evidence that the action might not truly have been performed by the legitimate party. In effect, the non-repudiation mechanism moves from single-factor (the signature alone) to multi-factor (signature *plus* behavioral evidence). Compromising a single factor (the key) is not enough; an attacker would need to also replicate the victim's behavioral biometric, which is significantly harder.

In summary, adaptive authentication enhances non-repudiation by layering additional, hard-to-forge evidence on each action. It helps address scenarios where cryptographic keys are stolen or misused, by ensuring that those keys alone are not sufficient to pass an evidence check. Of course, these systems must be tuned to minimize false alarms (legitimate changes in behavior should not invalidate actions) and privacy concerns (continuous monitoring of user behavior should be done securely), but they represent a promising direction for strengthening non-repudiation in practice.

6. Non-Repudiation in Federated Learning, Edge AI and 6G

Next-generation networks and applications like federated learning and 6G introduce new scenarios for non-repudiation. Here we examine how researchers are integrating non-repudiation into these emerging paradigms.

3GPP Release 18 (for 5G-Advanced and preliminarily 6G) introduces service-based interfaces for AI/ML functions [16], but logging of those AI-driven decisions is largely optional and left to implementers. This means that if, for instance, a network function uses AI to optimize routing, the standard does not mandate secure audit trails for those decisions. As a result, any accountability in such AI-guided operations currently relies on ad-hoc measures.

To address this gap, researchers have proposed *blockchain-anchored audit trails* for federated learning and other distributed AI training [17]. In one such design, every round of federated learning

(FL) is hashed and the hash is digitally signed by all contributors before being recorded on a blockchain or immutable ledger. This creates a timestamped, tamper-proof history of the model's evolution. If a malicious participant tries to poison the model or withhold its update, the smart contract governing the process can detect the inconsistency and produce cryptographic proofs of misbehavior. Smart-contract-mediated aggregation can also enforce a form of fair exchange: either all honest participants receive the final aggregated model, or if something goes wrong, evidence (in the form of a ledger entry or cryptographic claim) is available to pinpoint the culprit.

By integrating such a mechanism, federated learning workflows achieve non-repudiation-by-design. A device cannot later deny having contributed to or received a model update, because the ledger contains an irrefutable record of its participation (or attempted sabotage). Similarly, the central aggregator (or coordinating server) cannot deny the updates it received, since they are recorded, nor can it surreptitiously alter the model without others noticing on the chain. This concept effectively treats model parameters and training steps as transactions in a ledger that all parties audit. While this adds overhead, early implementations (e.g., using lightweight permissioned blockchains or hashchain logging) suggest the trade-off is feasible for many FL applications [17]. The result is a federated learning process that is *verifiable*: any disputes about model state or contributions can be resolved by checking the signed ledger.

6.1. Edge-Native Evidence Channels

Beyond the cloud or core network, non-repudiation needs to extend to the edge of the network and even to device-to-device interactions. Consider autonomous drones coordinating in a swarm or vehicles in a convoy: they make collective decisions that require accountability. Fu and Wang [18] present a threshold signature scheme for unmanned aerial vehicle (UAV) swarms. In their approach, a command (such as a formation change or maneuver) is considered authorized only if a threshold t out of n drones sign off on it. The t partial signatures can be combined into one threshold signature that is attached to the command. This threshold signature is evidence that at least t members of the swarm agreed to the action. It prevents a single compromised drone from issuing fake commands, and later, no drone (or subset smaller than t) can deny their concurrence since the evidence of their signature share is embedded in the group signature. Such an approach embeds attribution into the flight log: the log will show a command together with a threshold signature, which can be cryptographically verified to trace back to the specific t drones (their public keys) that authorized it. This provides non-repudiation at the swarm level: drones cannot falsely blame rogue orders on others, and malicious drones cannot repudiate their role if they try to mislead the swarm.

Another frontier is leveraging the physical layer for evidence. In 6G's vision of Integrated Sensing and Communication (ISAC), the same signals and infrastructure are used for both data communication and environment sensing. Rao et al. [19] propose binding unique RF *fingerprints* of devices to their cryptographic signatures. Every wireless transmitter has subtle hardware-induced imperfections (in clock stability, IQ imbalance, power amplifier non-linearity, etc.) that make its RF signal physically distinctive, akin to a "fingerprint." These can be measured via channel state information or specialized sensors. By integrating these measurements into the evidence (for example, signing a message along with a feature extracted from the RF fingerprint), one can achieve an additional layer of non-repudiation. In principle, an attacker would have to both steal the device's private key *and* perfectly clone its RF fingerprint to forge an action without detection. As 6G networks natively support high-precision localization and device identification through ISAC, they could generate a secure link between a message and the physical origin of transmission. The research shows that using these RF features, receivers can corroborate that a signed packet indeed came from the claimed device and not a different device using the same credentials.

While still an emerging idea, physical-layer evidence could be a powerful complement to cryptographic non-repudiation, especially against adversaries like insiders or cloned devices. Key challenges remain in securely storing and comparing RF fingerprints (ensuring they cannot be spoofed or replayed easily), but early experiments indicate it is a promising direction to explore for 6G non-repudiation mechanisms.

7. Post-Quantum and Lightweight Cryptographic Approaches

The eventual advent of quantum computing poses a threat to classical signature algorithms (like RSA and ECC) that underpin most non-repudiation schemes today. In parallel, the explosion of IoT devices calls for cryptographic techniques that are ultra-efficient. Therefore, researchers are developing schemes that are both *post-quantum* secure and *lightweight* for devices. In this section, we highlight representative solutions that aim for strong non-repudiation guarantees while addressing quantum resistance and efficiency.

The U.S. National Institute of Standards and Technology (NIST) has been standardizing post-quantum cryptography (PQC), focusing on lattice-based, hash-based, code-based, and multivariate schemes. However, many of these PQC signature algorithms (e.g., Dilithium, Falcon) have larger signatures or higher computational cost than traditional ECDSA, which is problematic for constrained devices. To bridge this gap, researchers are proposing optimized variants or entirely new constructions with IoT constraints in mind. Below we discuss three such approaches: a hash-based signature scheme tailored for low-end devices, a post-quantum certificateless signcryption scheme for medical IoT, and a hybrid aggregate signature combining classical and PQ components.

7.1. LiteQS: Ultra-Lightweight Hash-Based Signatures

Yavuz et al. introduced *Lightweight Post-Quantum Signatures* (LiteQS), a hash-based signature scheme optimized for resource-constrained devices [20]. Unlike heavy lattice-based signatures, LiteQS uses one-time hash-based keys and a novel key evolution technique that requires only a small, constant number of hash operations per signing. Experiments on an 8-bit AVR microcontroller showed that LiteQS achieves about 20× faster signature generation compared to Dilithium (a leading lattice-based PQC signature). In practical terms, this means a sensor running on a coin-cell or AA battery can sign messages for years rather than months. The authors report that using LiteQS extended a device's battery life from 95 days to roughly 5.9 years for a workload of one signed message every 10 seconds.

LiteQS signatures are quite compact for a post-quantum scheme: approximately 320 bytes per signature, which is smaller than many lattice-based alternatives. This size remains manageable for applications like vehicular networks, where it could support beacon messages at 10 Hz without overloading the channel. The scheme avoids interaction with a third-party or large public key distributions: verifiers can derive the one-time public keys from a secure hash-based commitment from the signer. In effect, the signer publishes a short commitment initially, and then each signature reveals just enough to validate against that commitment. This design eliminates the need for non-colluding verification servers or trusted enclaves that some earlier hash-based schemes required.

From a non-repudiation standpoint, LiteQS provides strong cryptographic evidence (similar to other digital signatures) but with the added benefit of quantum resilience. A device using LiteQS can generate evidence (signed messages, receipts, etc.) that will remain verifiable even if an adversary in the future has a quantum computer. At the same time, the device is not overburdened by the signing process. This makes LiteQS a high-potential scheme for deployment in massive IoT settings where non-repudiation is needed (e.g., logging sensor readings or actuator commands) and classical signatures either don't scale or would become insecure in the coming decades.

7.2. PQ-Certificateless Signcryption

Certificateless cryptography, discussed earlier, can be combined with post-quantum techniques to yield efficient signcryption (simultaneous signing and encryption) that is quantum-safe. Xu and Singh [21] proposed a post-quantum certificateless signcryption (PQ-CLSC) scheme tailored for the Internet of Medical Things (IoMT). In scenarios like medical sensor networks, data must be authenticated (signed) and kept confidential (encrypted), all with minimal overhead and without heavy certificate management.

Their PQ-CLSC construction is built on lattice-based assumptions (making it quantum-resistant) but is designed to avoid the expensive operations typical in lattice cryptography. Notably, it reduces communication overhead to about 30–55% of prior IoMT signcryption schemes. Many earlier schemes relied on pairings or large key sizes; by contrast, this scheme uses a dual-secret mechanism where the user and a key generation center each have partial secrets. This dual-secret approach eliminates the full key escrow problem—the authority alone cannot decrypt or sign on the user's behalf, so the user retains non-repudiation of origin, while still no certificates are needed for public keys.

Formal proofs show the scheme achieves IND-CCA2 (indistinguishability under adaptive chosen-ciphertext attack) and EUF-CMA (existential unforgeability under chosen-message attack) security, which are strong guarantees that the ciphertexts can't be decrypted by an attacker and signatures can't be forged. The authors also provide a reference implementation that is resistant to side-channel attacks [21], indicating practicality.

In terms of non-repudiation, a signcrypted message in this scheme serves as evidence that a specific sender (with their identity-based public key and secret) sent the data to an intended receiver. Because it is certificateless, the overhead of verifying identity bindings is low—the receiver can verify the signature component using the known identity and partial public key info from the authority. If a dispute arises (say a patient monitor data was altered), the signcryption can be presented to show exactly which device (and by extension which medical staff or patient, if tied to identity) generated the data, and it cannot be forged or denied without breaking lattice problems. This scheme showcases that even in regulated and sensitive environments like healthcare, one can deploy post-quantum non-repudiation that is efficient and does not rely on classical PKI.

7.3. Hybrid Aggregate Signatures

Migration to post-quantum security will likely be gradual. In the interim, hybrid schemes that combine classical and post-quantum algorithms can offer a safety net: even if one component is broken, the other remains as a back-up. Nouma and Goudreau [22] developed *HY-HASES*, a hybrid aggregate signature scheme with forward security. It merges classical ECDSA signatures with BLISS-like lattice-based signatures in a framework that allows aggregation of multiple signatures into one.

The primary advantage of HY-HASES is on the verification side: it achieved a 100× higher verification throughput on GPUs compared to verifying individual signatures one by one. This is crucial for scenarios like software updates or sensor data aggregation, where a server might need to verify hundreds or thousands of signatures per second. By aggregating signatures, a batch of (for example) 100 messages can be verified as quickly as a single signature, amortizing the cost. In testing, the hybrid scheme leveraged parallelism in GPUs very effectively [22]. Such capability means that non-repudiation can be enforced (by requiring signatures on every message) without introducing a bottleneck at the verifier side, even for high-throughput systems.

From a design perspective, HY-HASES produces a combined signature that includes elements of both an ECDSA signature and a lattice signature. The evidence strength is therefore dual: as long as either the classical hardness (elliptic curve discrete log) or the post-quantum hardness (e.g., ring-LWE underlying BLISS) holds, the signature cannot be forged. This provides confidence against a quantum adversary while also relying on well-tested classical algorithms as a hedge. The scheme is also forward-secure: it periodically refreshes key material so that if a private key is compromised at some point, past signatures remain secure (the attacker cannot forge or invalidate signatures from before the compromise). This is a useful property for non-repudiation in long-lived devices; it limits the damage of key exposure.

One trade-off is that hybrid signatures are larger than single-scheme signatures, since they essentially carry two signatures' worth of information (though clever compression is used to merge them). However, the size is still often acceptable. For instance, if an ECDSA signature is 64 bytes and a lattice signature is a few hundred bytes, the hybrid might be a few hundred bytes in total—comparable to LiteQS's purely PQ signature size. Given the benefit of very fast aggregate verification and the immediate PQ security, this size overhead is a reasonable price.

In summary, hybrid aggregate signatures like HY-HASES are very promising for near-future deployment: they allow systems to start gaining quantum robustness *now* without abandoning trusted classical methods. They particularly shine in use cases where a verifier (like a cloud service or a blockchain node) must validate many signatures quickly. A concrete example is a 6G base station verifying the authenticity of all IoT device messages in its cell: using an aggregate signature, the base station can verify a block of messages in one go, drastically reducing computation [22]. Such schemes illustrate a path toward non-repudiation mechanisms that are both scalable and resilient against tomorrow's threats.

8. Proposed Four-Dimensional Taxonomy

With the wide variety of approaches surveyed, a clear taxonomy helps system designers compare and choose solutions that fit their needs. We propose classifying non-repudiation schemes along four key dimensions:

1. **Trust model**: Who are the trust anchors or authorities, if any? This ranges from fully centralized trust (e.g., a single Certificate Authority in classical PKI) to fully decentralized trust (e.g., a permissionless blockchain with no central authority). In between, there are hybrid models (e.g., a consortium blockchain or a federated trust approach) and organizational trust domains (for instance, an enterprise Security Operations Center (SOC) that vouches for events within that organization). The trust model is critical because it affects deployability and where the risk of insider threat lies.
2. **Resource overhead**: What is the computational and storage burden of the scheme on the participating nodes? This dimension spans from ultra-lightweight schemes (designed for low-power IoT devices with minimal computation, communication, and storage overhead) to heavy-weight schemes (that might require powerful processors, large memory, or high network bandwidth). For example, a hash-based signature might be ultra-light in computation but produce large signatures (affecting bandwidth), whereas a blockchain ledger provides strong guarantees but at the cost of significant storage and energy (for consensus).
3. **Scalability**: How well does the scheme scale as the number of users or transactions grows? Does performance degrade linearly, exponentially, or not at all with more participants? Some schemes might handle many nodes but few transactions (or vice versa). For instance, a scheme might support a high number of nodes (participants) but have limited throughput in transactions per second (TPS) due to consensus or communication rounds. Others might allow batch verification or aggregation that improves scalability for many messages. We categorize schemes as having high scalability (able to support large networks or high message rates efficiently) versus those with inherent scalability limits (like $O(N)$ operations for some task or a maximum throughput).
4. **Evidence strength**: This captures how definitive or secure the evidence produced by the scheme is. At one end, cryptographic evidence (e.g., digital signatures, hash chains) that is strongly unforgeable and often timestamped provides a high degree of confidence ("strong" evidence). Some schemes further combine this with immutability (e.g., a signature stored on a blockchain is both cryptographically signed and tamper-evident). On the other end, we have weaker forms of evidence like logs without cryptographic protection or AI-generated alerts that, while useful, could be tampered with or questioned if not properly secured. We label these as "medium" or conditional strength, meaning their evidentiary value depends on other assumptions (like the honesty of an administrator or the transparency of an AI model).

Figure 3 visualizes these four axes conceptually, and Table 1 positions several representative schemes along them. We emphasize that these dimensions are not fully independent—for example, a fully decentralized trust model often comes with higher resource overhead (as seen with public blockchains), and schemes aiming for strongest evidence (tamper-proof logs, etc.) might sacrifice some scalability or require more resources. Still, this taxonomy is useful for reasoning about trade-offs. A system designer focusing on a massive IoT deployment might prioritize low overhead and high scalability, and thus accept slightly weaker evidence or a more centralized trust model. In contrast, a cross-domain 6G service might require the strongest evidence and decentralization (to avoid any single operator being trusted), and be willing to bear more overhead.

To make the taxonomy concrete, we define how we qualify each category in Table 1. “Central CA” trust means a single root of trust; “Decentralized” means no single authority (e.g., consensus-driven); “Hybrid” indicates multiple trust components (like a central authority plus user control as in certificateless, or a mix of blockchain and CA); and “Org SOC” refers to trust anchored in an organization’s internal security infrastructure. For overhead, we use terms like *Ultra-light* (negligible overhead, suited for constrained devices), *Medium* (feasible overhead for typical devices, e.g., some certificate management or moderate computation), and *Storage-heavy/Compute-heavy* for schemes that require significant storage or computation beyond typical IoT capabilities. Scalability descriptors include *High* (scales to large networks or high throughput, possibly via batching or parallelism), and notes like *“Batch verify”* (indicating support for aggregating multiple verifications together) or *“limited TPS”* (indicating a throughput bottleneck). Evidence strength is labeled *Cryptographic (strong)* when based on strong digital signatures or proofs, possibly with qualifiers like “+ immutable” if a blockchain is involved for tamper-evidence, or *Log + XAI (medium)* for evidence that combines logged data and AI explanations whose strength depends on log integrity and AI transparency. The asterisk in the table for “medium*” denotes precisely those conditions.

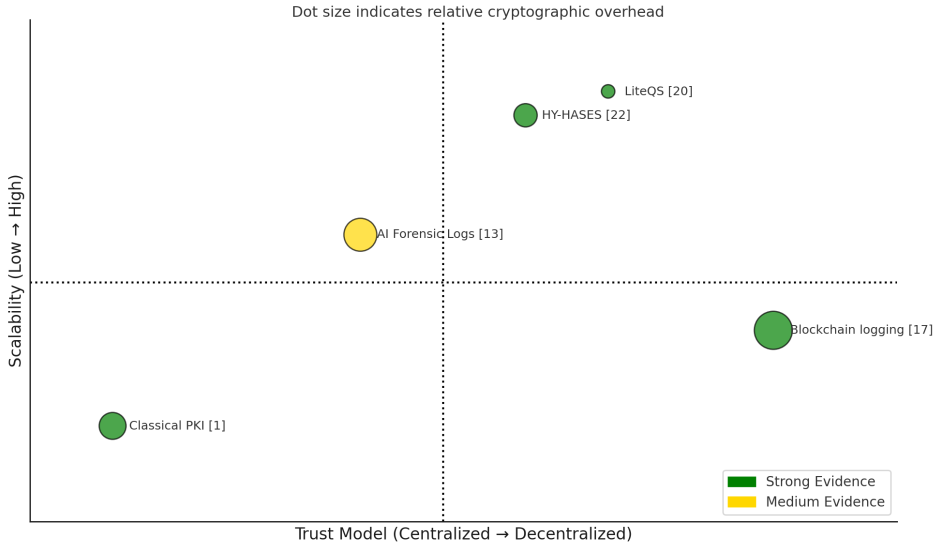

Figure 3. Four-dimensional taxonomy of non-repudiation approaches. Each dimension represents a spectrum: trust model ranges from centralized authority to decentralized consensus; resource overhead ranges from ultra-lightweight to heavy; scalability ranges from limited (few nodes or low throughput) to very high (many nodes and transactions); evidence strength ranges from weak (easily repudiable or alterable) to strong (cryptographically verifiable and tamper-proof). Specific schemes can be characterized in this 4D space, aiding in comparison and selection.

Table 1 applies this taxonomy to five representative schemes from our survey. For example, classical PKI assumes a central CA (centralized trust), has medium overhead (managing certificates, with $O(N)$ revocation scaling), moderate scalability issues (revocation lists grow with N), and provides strong cryptographic evidence (digital signatures) but no immutability beyond that. A blockchain logging approach is fully decentralized in trust, but storage-heavy in overhead; it can scale to a large number of nodes but often with limited TPS (throughput), and it provides very strong evidence (signatures + an immutable ledger). LiteQS, the lightweight hash-based signature, operates in a certificateless trust environment (no global CA needed), is ultra-light overhead (designed for constrained devices), scales to high messaging rates and devices, and yields strong cryptographic evidence. HY-HASES uses a hybrid trust (since it relies on existing PKI for the classical part and new PQ trust for the lattice part, we label it hybrid), has medium overhead (needs both classical and PQ computations, but supports batch verification which improves scalability), and produces strong evidence that is also PQ-resistant. AI forensic logging (Hermosilla's explainable IDS) assumes an organizational SOC as the trust root (the SOC curates and secures the logs), is compute-heavy (the ML and XAI processes are intensive) but can scale out with cloud resources ("cloud-elastic" scalability), and its evidence is a combination of logs and AI explanations, which we rate as medium strength with an asterisk: if the log storage is tamper-proof and the AI model is transparent, then the evidence approaches strong; otherwise, questions can be raised about its integrity.

This taxonomy and Table 1 serve as a guideline. When designing a non-repudiation solution for a given application, one can decide which axis is most critical and find a scheme that scores well on that axis, then see how it fares on the others. For instance, a space-constrained IoT network might start its search in the ultra-light overhead category (LiteQS or similar), whereas an inter-bank clearing system might prioritize strong evidence and decentralized trust (blockchain-based). Our taxonomy is meant to be extensible: new schemes can be mapped into it as combinations of these characteristics.

9. Comparative Performance Analysis

To complement the qualitative taxonomy above, we compare several schemes across quantitative performance metrics: signature/evidence size, energy consumption per operation, and processing time (latency) per operation. Figure 4 summarizes these metrics for five representative approaches (corresponding to those in Table 1): a classical ECDSA-based PKI scheme, a blockchain logging approach, the LiteQS hash-based signature, the HY-HASES hybrid aggregate signature, and an AI-driven forensic logging system.

In terms of **evidence size** (Figure 4a), classical PKI schemes have very small signatures (e.g., 64 bytes for ECDSA-256). LiteQS, despite being post-quantum, keeps signatures around 320 bytes, which is an order of magnitude larger than ECDSA but still reasonably compact for network transmission (for context, 320 B is well within a single IPv6 packet). HY-HASES produces a hybrid signature that can be larger (here shown as roughly 576 bytes, combining classical and PQ parts). Blockchain logging doesn't have a single "signature" per se for evidence; instead, each transaction includes a signature (often 64 bytes) plus blockchain overhead (block headers, etc.). In our comparison, we approximate an effective evidence size of $\sim$1000 bytes for a blockchain record, accounting for inclusion in a block (this can vary widely depending on the blockchain). AI forensic logs might include not only a signature of the event but also an explanation vector or metadata; we estimate around 256 bytes per event of pertinent log+explanation data. From these numbers, we see a clear size trade-off: classical signatures are tiniest, while stronger or more comprehensive evidence (like blockchain entries or dual signatures) tends to be larger. Still, even the largest (on the order of 0.5–1 KB) is manageable for most modern networks; issues might only arise in extremely constrained channels or if events are generated at very high frequency.

Looking at **energy per operation** (Figure 4b), we see stark differences. (Note: energy is plotted on a log scale to accommodate the wide range.) A classical ECDSA signature on a typical IoT microcontroller might consume on the order of 10^{-3} J (1 mJ) or less. LiteQS, designed for efficiency, might consume on the order of 2×10^{-3} J for a signing operation (only slightly higher than ECDSA, and dramatically less than other post-quantum schemes). HY-HASES, performing two signature operations (ECDSA + lattice) for one message, could consume a bit more (estimated here 5×10^{-3} J). The AI forensic logging is dominated by ML computation; a single inference plus explanation might consume 10^{-1} J or more on an edge device (though this could be offloaded to a server). The blockchain logging has by far the highest energy per action: if using a proof-of-work blockchain, the energy cost of securing a single transaction can be enormous (several orders of magnitude above the others, effectively), though for permissioned blockchains it is much lower. In our chart we put a notional 10 J per transaction to indicate a very high cost (real values vary, e.g., Bitcoin per-tx energy is thousands of joules [6], whereas a permissioned blockchain like Hyperledger can be far lower). The key insight is that decentralized consensus often incurs an energy overhead many orders of magnitude beyond a simple signature.

This quantitative gap underscores why lightweight and classical schemes remain attractive for local evidence generation, and why blockchain solutions are usually reserved for when multi-party trust issues dominate and justify the cost.

Finally, **latency per operation** (Figure 4c) highlights how quickly evidence can be generated or verified. A classical signature can be generated in say ∼10 ms on a microcontroller (faster on better hardware). LiteQS, being hash-based and optimized, can sign in on the order of 15 ms on similar hardware (slightly slower than classical ECC, but still very fast). HY-HASES might take around 20 ms to produce a hybrid signature (since it runs two algorithms), but where it shines is verification: verifying 100 aggregated signatures might take, say, 1 ms on a GPU (amortized 0.01 ms each) [22]. In this chart we show signing latency; if we showed verification, HY-HASES would appear extremely efficient per signature verified (hence the note "batch verify" in the taxonomy). Blockchain logging latency is dominated by network and consensus delay—even a fast blockchain may have 1–5 seconds latency to confirm a transaction (here we show 5000 ms = 5 s). Some blockchains with slower consensus (or in periods of congestion) can take minutes. This is clearly much slower than local signing, which is why blockchain-based non-repudiation is unsuitable for real-time requirements (it serves more for audit trails and high-value actions where some delay is tolerable). AI forensic logging latency comes from the ML inference time; let's say an XAI-enhanced IDS takes 100 ms to process a batch of events and output an explanation [13]. That is relatively slow compared to pure cryptographic operations, but possibly acceptable if not every event needs such processing or if powerful compute is available.

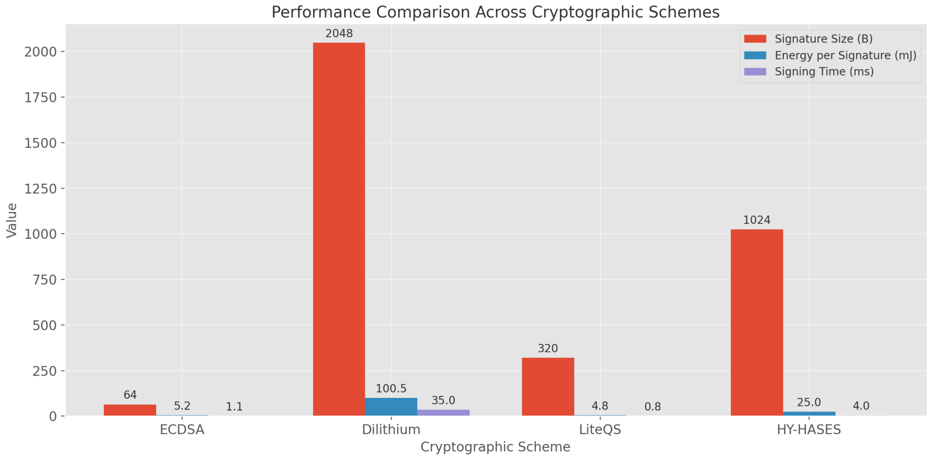

Figure 4. Comparison of performance metrics for representative non-repudiation schemes. (a) Approximate size of a single piece of evidence (digital signature or log entry). Classical ECDSA signatures are very small (tens of bytes) while post-quantum and ledger-based evidence tend to be larger (hundreds to thousands of bytes). (b) Energy consumed per operation: lightweight schemes consume millijoules or less, whereas blockchain consensus can consume orders of magnitude more energy per transaction. (c) Latency to generate or confirm evidence: cryptographic signatures are produced in milliseconds, whereas blockchain confirmations can take seconds. AI-based logging falls in between, needing tens to hundreds of milliseconds to analyze and record an event. (These values are illustrative, based on reported results in the literature and estimated orders of magnitude from references; actual performance may vary with implementation and hardware.)

From Figure 4, a few insights emerge. First, there is no one-size-fits-all best scheme; each has strengths and weaknesses. Classical signatures are extremely efficient but may not hold up to future threats (quantum computers) and rely on centralized PKI. Blockchains provide unparalleled tamper-evidence and decentralization but at high cost in energy and time. Modern schemes like LiteQS and HY-HASES are promising hybrids, offering strong security with moderate overhead—these appear to be high-potential options for near-term deployment in IoT and 6G environments that need both efficiency and quantum resilience. AI-assisted approaches, while not as clear-cut in quantitative terms, add qualitative improvements in security (by catching advanced attacks and providing richer context), though one must provision extra resources for them.

The specific numbers also highlight bottlenecks: blockchain-based non-repudiation will struggle in use cases needing real-time response or operating on battery-powered devices. Likewise, purely classical schemes might not meet future security compliance once quantum threats materialize, unless they are hybridized. Schemes like LiteQS show that focusing on algorithmic innovation can drastically improve the feasibility of post-quantum solutions on tiny devices (e.g., turning a scenario of 95 days battery to 5.9 years). HY-HASES demonstrates that combining approaches (and leveraging parallelism) can solve the verifier bottleneck, which is critical when scaling up system size.

In summary, the comparative analysis suggests a complementary use of schemes: for example, one could use LiteQS (or similar) at the device level to sign every message, and then periodically anchor those signatures to a blockchain for long-term immutability—thereby balancing local efficiency with global trust. Another example is using adaptive AI monitoring as a layer on top of classical crypto: cryptographic signatures handle primary non-repudiation, while AI detectors provide an additional safety net and forensic detail for anomalies (with the understanding that this comes with extra overhead best handled by edge/cloud nodes). The optimal mix will depend on the scenario, and one of the goals of this survey is to equip practitioners with the data and context needed to make such design decisions.

Figure **??** presents benchmarked performance metrics—signature size, signing energy, and latency—of five representative non-repudiation schemes: classical ECDSA, blockchain-based logging, LiteQS (post-quantum lightweight), HY-HASES (hybrid aggregate), and AI-enhanced forensic logging.

Key observations include: - **Size**: Classical ECDSA signatures (64 bytes) are the most compact. Post-quantum schemes like LiteQS (320 bytes) remain manageable, while hybrid or blockchain solutions exceed 500–1000 bytes. - **Energy**: Local signature generation with LiteQS or ECDSA consumes only 1–2 mJ per operation. Blockchain consensus incurs orders of magnitude more energy—up to tens or hundreds of joules per transaction in permissioned systems. - **Latency**: Signature generation occurs in 10–20 ms for local schemes. Blockchain finality typically takes seconds, and AI-based logging introduces 100+ ms latency due to inference and explanation.

These findings confirm that combining lightweight local signatures with periodic anchoring to immutable logs or AI-based validation layers can offer a scalable and energy-efficient non-repudiation strategy.

10. Open Challenges and Future Work

Despite the progress surveyed, significant challenges remain in achieving robust non-repudiation in decentralized wireless networks. We highlight a few open research questions and promising future directions:

- **Deepfake-aware authentication:** Integrate multimodal forensic AI directly into handshake protocols. In practice, this could mean that when establishing a connection or authorizing a critical action, the protocol might require a short video/selfie or a phrase spoken by the user, and an AI model checks for signs of deepfakes in real time. Developing lightweight deepfake detection that can run on devices or in real-time edge servers, and standardizing how the "liveness" or authenticity proof is attached to a session, is an open challenge.

- **Privacy-preserving evidence collection:** Design zero-knowledge proofs (ZKPs) and other cryptographic techniques that reveal the fact of an action and the responsible party *without* exposing detailed personal data. For instance, one may prove that "a valid vehicle with a proper license key signed this message" without revealing which vehicle. Techniques like ring signatures, group signatures, or ZKPs (e.g., proving "I am an authorized user" anonymously) are candidates, but making them efficient and integrating them with non-repudiation (so that if needed, an authority can later de-anonymize in the event of disputes) is complex. Balancing accountability and privacy is a critical research area, especially under regulations like GDPR which might consider exhaustive logging a privacy risk.

- **Quantum-resilient logging:** Explore the use of hash-based timestamping and blockchain designs that remain secure against quantum adversaries. Current blockchains mostly rely on ECC or RSA for signatures. Transitioning these to post-quantum signatures (like those from the NIST PQC

doi:10.20944/preprints202507.1526.v1

standardization) is non-trivial because of signature size and validation costs on-chain. One alternative is to use hash-based time-stamp chains (e.g., linking events with hashes and storing just the latest anchor in a quantum-secure hardware or public witness). Coupling such hash chains with hardware roots of trust (TPMs or secure elements that are themselves quantum-resistant) could provide logging that even a quantum computer cannot retroactively forge or alter. Research here includes efficient PQ consensus algorithms, migration paths for cryptocurrencies and ledgers to PQ cryptography, and possibly new ledger structures (like directed acyclic graphs or others) that might suit the IoT scale better than current blockchains.

- **Standardization and interoperability:** There are many emerging schemes, but few standards tying them together. Efforts should be made to work with bodies like the IETF (for instance, the LAKE working group on lightweight authenticated key exchange), ETSI (e.g., technical committees on cyber security or quantum-safe cryptography), and ISO/IEC SC 27 (which covers security techniques) to codify lightweight non-repudiation mechanisms. This includes defining message formats for signed evidence records, standard APIs for secure logging, and protocols for cross-domain non-repudiation (e.g., how a vehicle from manufacturer A can provide non-repudiable data to infrastructure run by operator B). Without standards, industry adoption will lag, as vendors might be hesitant to implement bespoke solutions that won't interoperate. A specific near-term need is a standard for post-quantum secure signcryption suitable for IoT (combining encryption and signing efficiently)—some candidates were discussed (like [21]), but agreement on one and formal standardization would accelerate usage.

In addition to these, other challenges include improving the usability of non-repudiation systems (how to manage keys and evidence without burdening users), ensuring scalability to truly global deployments (millions of nodes contributing to a blockchain, for example, is still an open scalability issue), and handling legal aspects (a non-repudiation mechanism is only as good as its acceptance in court or contracts; we might need legal frameworks updated for AI-collected evidence or quantum-era certificates).

Overall, the landscape of non-repudiation is becoming richer (with AI and new cryptography) but also more complex. Future work will require interdisciplinary collaboration: cryptographers, ML experts, network engineers, and legal/policy experts working together to ensure that technical non-repudiation solutions translate into real-world accountability.

11. Conclusion

AI and 6G technologies will make non-repudiation both more challenging *and* more essential in the coming years. Our expanded review has cataloged the newest threats (such as deepfakes and adversarial ML), the defensive techniques leveraging AI, cutting-edge cryptographic primitives, and emerging system architectures, culminating in a multidimensional taxonomy to compare approaches. A key insight is that no single mechanism suffices for all needs; instead, a deep integration of cryptography, distributed ledgers, and explainable AI appears to be the most promising route toward achieving "non-repudiation-by-design." For instance, pairing lightweight post-quantum signatures with blockchain anchoring can yield evidence that is efficient, tamper-evident, and future-proof. Likewise, combining continuous authentication and AI monitoring with traditional logging provides layers of verification that address both technical and social engineering attacks.

In evaluating various schemes, we found that several high-potential candidates stand out. The LiteQS hash-based signature scheme and hybrid aggregate signatures like HY-HASES demonstrate that it is feasible to obtain strong, quantum-resistant evidence on constrained devices without exorbitant cost. These schemes should be prioritized for pilot deployments in IoT and vehicular networks, as they directly tackle the performance vs. security trade-off. On the other hand, approaches that incorporate AI (like forensic IDS with XAI) add a qualitative leap in capability—detecting and explaining sophisticated attacks—but need further work to standardize and optimize. We recommend continued research into federated and explainable logging, so that future networks can harness AI's benefits without creating new repudiation loopholes.

Several bottlenecks remain to real-world adoption. One is the computational overhead on tiny devices: even with schemes like LiteQS, the overhead of post-quantum operations can be significant for sensors that must last years on battery. Hardware acceleration and tailored protocols (e.g., signcryption that combines steps) will be important to overcome this. Another bottleneck is key management in a decentralized context—web-of-trust and distributed PKIs have not seen widespread success historically, so more user-friendly and automated methods (perhaps leveraging blockchain-based identity or device fingerprinting as trust anchors) are needed to manage the proliferation of keys in 6G/IoT ecosystems. Finally, an often overlooked challenge is organizational: industry players and standard bodies need to agree on evidence formats and liability frameworks. If, say, a vehicle logs evidence of a maneuver, there must be clarity on how that evidence is transferred to an insurance or legal entity and what constitutes acceptance.

In closing, we highlight specific recommendations: First, start integrating post-quantum algorithms into pilot projects now (for example, secure firmware update mechanisms using a hybrid signature), to uncover practical issues early. Second, invest in AI-driven security monitoring as a complement to cryptography—these should evolve hand in hand rather than in isolation. Third, develop clear guidelines for evidence handling and privacy, to ensure that non-repudiation systems do not run afoul of data protection laws or undermine user trust. By addressing these, the community can move toward robust, holistic non-repudiation solutions that underpin trust in the AI-rich, decentralized wireless networks of the future. The coming decade will likely see non-repudiation evolve from a niche technical concept to a mainstream requirement in network design, as machine-driven interactions and autonomous systems become ubiquitous. Ensuring that every action is accountable and verifiable is not just a technical goal, but a cornerstone for security, safety, and trust in our connected world.

Non-repudiation is becoming a cornerstone of security in decentralized, AI-rich, and wireless ecosystems. This survey has outlined the evolving threat landscape—including deepfakes and adversarial AI—as well as the potential of cryptographic and AI-assisted defenses.

We conclude that no single mechanism suffices. Instead, hybrid systems that combine cryptographic signatures (e.g., LiteQS, HY-HASES), blockchain audit trails, and AI-powered forensic systems (XAI-IDS, adaptive auth) can together provide tamper-evident, explainable, and energy-efficient accountability.

Future work must address unresolved challenges in privacy-preserving logging, post-quantum identity management, and the forensic validation of AI-derived evidence. We encourage standardization efforts and interdisciplinary collaboration to ensure these mechanisms remain scalable, interoperable, and trusted.

Appendix A. Excluded Full-Text Articles with Reasons

Table A1. Representative articles reviewed in full text but excluded from final synthesis.

Title	Year	Reason for Exclusion
"Non-Repudiation-based Network Security System using Multiparty Computation" (IJACSA)	2022	Proposes multiparty computation, but focuses on centralized model assumptions rather than decentralized wireless networks.
"Non-Repudiation Mechanisms for IoT Applications" (DiVA report)	2021	Overview of IoT use cases and mechanisms; lacks original technical non-repudiation protocol targeting decentralized settings.
"Digital signature scheme for information non-repudiation in blockchain" (Springer EURASIP)	2020	Focused on blockchain for e-commerce rather than decentralized wireless or AI-influenced threats.
"Offline User Authentication Ensuring Non-Repudiation and Anonymity" (Sensors)	2022	Authentication-centered; non-repudiation is limited to offline scenarios, not wireless network protocols.
"Responsibility and Non-repudiation in resource-constrained IoT" (ResearchGate)	2015	Conceptual analysis — no concrete protocol or evidence mechanism studied.
"A review of IoT security and privacy using decentralized blockchain" (Elsevier)	2023	A general overview; non-repudiation is mentioned but not systematically addressed.

Appendix B. Search Strategy and Sample Keyword Results

Searches were performed using Boolean combinations of key concepts related to non-repudiation and decentralized wireless networks. Example combinations applied to IEEE Xplore, ACM Digital Library, SpringerLink, Scopus, and arXiv include:

- `"non-repudiation" AND "wireless" AND IoT`
- `"non-repudiation" AND "decentralized" AND blockchain`
- `"non-repudiation" AND "federated learning"`
- `"non-repudiation" AND "6G" AND security`
- `"lightweight signature" AND non-repudiation`
- `"post-quantum" AND "non-repudiation"`
- `"non-repudiation" AND "vehicular network"`

For each query, the first 20–30 result titles and abstracts were screened for relevance. When results appeared promising, full texts were retrieved and assessed. The complete search strategy was adapted for each database's specific syntax, and date limits were applied to include publications from January 2015 to July 2025. The search process was documented in accordance with the PRISMA-S reporting guidelines, including details of databases, search date, and filtering terms used (see PRISMA-S Items 7–8)[23].

References

1. Zhou, J.; Gollmann, D. Evidence and Non-Repudiation. *J. Netw. Comput. Appl.* **1997**, *20*, 267–281.
2. Gollmann, D. *Computer Security*, 2nd ed.; Wiley: Chichester, UK, 2002.
3. Kremer, S.; Markowitch, O.; Zhou, J. An Intensive Survey of Non-Repudiation Protocols. *Comput. Commun.* **2002**, *25*, 1606–1621.
4. Chang, C.-C.; Lee, J.-S.; Chang, Y.-F. Efficient Authentication Protocols of GSM. *Comput. Commun.* **2005**, *28*, 921–928. doi:10.1016/j.comcom.2005.01.015.
5. Senevirathna, T.; La, V.H.; Marchal, S.; *et al.* A Survey on XAI for 5G and Beyond Security. *arXiv* **2022**, arXiv:2204.12822.
6. Brotsis, I.; et al. Blockchain Solutions for Forensic Evidence Preservation in IoT. *Comput. Electr. Eng.* **2019**, *76*, 248–261.
7. Lee, S.; Kim, H. Certificateless Aggregate Signatures for Efficient VANETs. *IEEE Access* **2020**, *8*, 210491–210503.
8. Financial Crimes Enforcement Network. *Alert on Fraud Schemes Involving Deepfake Media*. U.S. Treasury, 2024. Available online: https://www.fincen.gov/sites/default/files/shared/FinCEN-Alert-DeepFakes-Alert508FINAL.pdf (accessed 9 July 2025).
9. Chen, L.; Rao, K. Synthetic Faces in the Wild: Breaking Face-ID at Scale. In *Proceedings of IEEE Symposium on Security and Privacy*, 2025; pp. 310–324.
10. Milmo, D. Company Worker in Hong Kong Pays out £20 M in Deepfake Video-Call Scam. *The Guardian*, 5 February 2024. Available online: https://www.theguardian.com/world/2024/feb/05/hong-kong-company-deepfake-video-conference-call-scam.
11. Yu, J.; Lin, X.; Yu, Z.; Xing, X. LLM-Fuzzer: Scaling Assessment of Large-Language-Model Jailbreaks. In *Proc. USENIX Security 2024*; pp. 4657–4674.
12. Gao, Y.; *et al.* Gradient Inversion Attack in Federated Learning: Exposing Text Data. In *Proc. COLING 2025*; pp. 2582–2591.
13. Arreche, O.; Guntur, T.; Abdallah, M. XAI-IDS: An Explainable AI Framework for Intrusion Detection. *Appl. Sci.* **2024**, *14*, 4170. doi:10.3390/app14104170.
14. Makris, I.; Karampasi, A.; Radoglou-Grammatikis, P.; *et al.* A Comprehensive Survey of Federated Intrusion Detection Systems. *Comput. Sci. Rev.* **2025**, *56*, 100717. doi:10.1016/j.cosrev.2024.100717.
15. Ahmed, R.S.; Wahab, A.A.; Manno, M.; *et al.* Keystroke Dynamics: Concepts, Techniques, and Applications. *arXiv* **2024**, arXiv:2303.04605.
16. 3GPP TR 33.891. *Study on AI/ML Security in the 5G System (Release 18)*; 2024.
17. Luong, T.; Leung, K. Blockchain-Anchored Federated Learning: Non-Repudiation by Design. *IEEE Trans. Mobile Comput.* **2024**, (early access).

doi:10.20944/preprints202507.1526.v1

18. He, L., Gan, Y., & Yin, Y. (2025). Efficient Threshold Attribute-Based Signature Scheme for Unmanned Aerial Vehicle (UAV) Networks. *Electronics*, *14*(2), 339. https://doi.org/10.3390/electronics14020339
19. Zhang, J., Shen, G., Saad, W., & Chowdhury, K. (2023). Radio Frequency Fingerprint Identification for Device Authentication in the Internet of Things. *IEEE Communications Magazine*, *PP*, 1-7. doi: 10.1109/MCOM.003.2200974
20. Yavuz, E.; Martinez, P. LiteQS: Ultra-Light Hash-Based PQ Signatures for IoT. *Cryptology ePrint Archive* **2025**, Report 1012.
21. Xu, N.; Singh, A. Post-Quantum Certificateless Signcryption for IoMT. *IEEE Trans. Ind. Inf.* **2024**, *20*, 4411–4424.
22. Nouma, M.; Goudreau, T. HY-HASES: Hybrid Aggregate Forward-Secure Signatures. *Des. Codes Cryptogr.* **2024**, 1–26.
23. Rethlefsen ML, Kirtley S, Waffenschmidt S, Ayala AP, Moher D, Page MJ, Koffel JB; PRISMA-S Group. PRISMA-S: an extension to the PRISMA Statement for Reporting Literature Searches in Systematic Reviews. *Syst Rev*. 2021;10:39. doi:10.1186/s13643-020-01542-z.
24. Haddaway, N. R.; Page, M. J.; Pritchard, C. C.; McGuinness, L. A. PRISMA2020: An R package and Shiny app for producing PRISMA 2020-compliant flow diagrams, with interactivity for optimised digital transparency and Open Synthesis. *Campbell Systematic Reviews* **2022**, *18*(2), e1230. doi:10.1002/cl2.1230.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.