International Journal of Scientific Research and Modern Technology (IJSRMT) ijsrmt.com
Volume 5, Issue 7, 2026
DOI: https://doi.org/10.38124/ijsrmt.v5i7.1551

Machine Learning-Based Detection of Data Privacy Vulnerabilities in Critical Smart Warehouse and Distribution Systems

Oluwadamilola Durowoju[1]; Isaac Quaye[2]

[1]Department of Computer Science Missouri University of Science and Technology, USA
[2]Department of Geography, Environment and Urban Studies Temple University, USA

Publication Date: 2026/07/13

Abstract

The rapid adoption of Internet of Things (IoT) devices, autonomous mobile robots, facial recognition surveillance systems, and cloud-integrated warehouse management platforms in U.S. smart warehouse and distribution facilities has created a complex and largely under-analyzed attack surface for data privacy violations. Worker biometrics, inventory records, supplier transactions, and operational telemetry are now generated and processed at scale across interconnected systems that span physical, network, application, and cloud layers. This article presents a comprehensive examination of machine learning-based techniques for detecting data privacy vulnerabilities in these environments, integrating findings from the latest research in deep learning intrusion detection, adversarial robustness, federated privacy-preserving learning, blockchain-secured IoT data sharing, and computer vision security. We analyze the threat landscape through five system layers, evaluate the performance of leading ML detection frameworks including WILS-TRS, TAD, FedPrIDS, and CNN-based vision audit systems, and examine their alignment with the U.S. regulatory landscape including CCPA, BIPA, NIST CSF 2.0, and Executive Order 14028. A phased implementation roadmap is proposed to guide warehouse operators from baseline intrusion detection to fully federated, blockchain-anchored privacy governance. Our analysis demonstrates that while no single ML paradigm addresses all vulnerability classes, a layered deployment combining federated learning, adversarial training, and blockchain data provenance offers the most robust and privacy-preserving solution for modern U.S. distribution infrastructure.

Keywords: *Machine Learning; Intrusion Detection; Data Privacy; Smart Warehouse; IoT Security; Federated Learning; Adversarial Attacks; CCPA; Deep Learning; Cybersecurity.*

I. INTRODUCTION

The warehousing and logistics sector is undergoing a profound technological transformation. Driven by e-commerce growth, labor market pressures, and supply chain resilience imperatives, leading distribution operators from Amazon fulfillment centers to third-party logistics (3PL) providers and cold-chain pharmaceutical distributors have deployed increasingly sophisticated smart warehouse systems. These environments integrate dense networks of IoT sensors, autonomous mobile robots (AMRs), barcode and RFID readers, computer vision surveillance, and cloud-connected warehouse management systems (WMS). The result is a highly automated, data-intensive operational environment that generates continuous streams of information about inventory movement, equipment status, and worker behavior.

This technological density, while operationally powerful, introduces significant and poorly understood data privacy vulnerabilities. Worker biometric data captured by facial recognition systems, precise location tracking data from wearable devices and AMR proximity sensors, and sensitive operational data transmitted across heterogeneous wireless protocols collectively constitute a vast repository of personally identifiable information (PII) that is continuously at risk of unauthorized access, exfiltration, or misuse. The intersection of operational technology (OT) and information technology (IT) in smart warehouses further expands the attack surface by eliminating the air-gap protections that historically

Durowoju, O., & Quaye, I. (2026). Machine Learning-Based Detection of Data Privacy Vulnerabilities in Critical Smart Warehouse and Distribution Systems. *International Journal of Scientific Research and Modern Technology*, *5*(7), 24–38. https://doi.org/10.38124/ijsrmt.v5i7.1551

separated industrial control systems from enterprise networks Shanthamallu et al. (2017).

Machine learning has emerged as the dominant paradigm for detecting privacy vulnerabilities and cybersecurity threats in complex IoT environments. Unlike rule-based systems that require manual signature updates, ML models can identify previously unseen attack patterns, adapt to evolving threat behaviors, and process the high-velocity, heterogeneous data streams characteristic of warehouse operations Alshuaibi et al. (2025). However, the deployment of ML-based detection systems in warehouse environments introduces its own privacy considerations: model training requires access to sensitive operational data, inference engines may inadvertently expose worker PII, and adversarial actors have demonstrated the ability to craft inputs that systematically evade even state-of-the-art detection systems Yin et al. (2019).

Recent research has also highlighted the increasing cybersecurity and operational risks associated with AI-enabled logistics infrastructure in the United States. Watara and Ahmed (2026) demonstrated how AI-driven security architectures can improve cyberattack detection and resilience across automated energy supply chains. This article addresses the intersection of these concerns by providing a comprehensive analysis of ML-based privacy vulnerability detection for U.S. smart warehouse and distribution systems. We examine the threat landscape across five system layers, evaluate leading ML detection frameworks, assess adversarial robustness challenges, compare centralized, on-device, and federated learning deployment architectures, and map technical capabilities against the evolving U.S. regulatory requirements. A phased implementation roadmap is presented to provide warehouse operators, security architects, and compliance officers with actionable guidance for deploying privacy-preserving ML detection at scale.

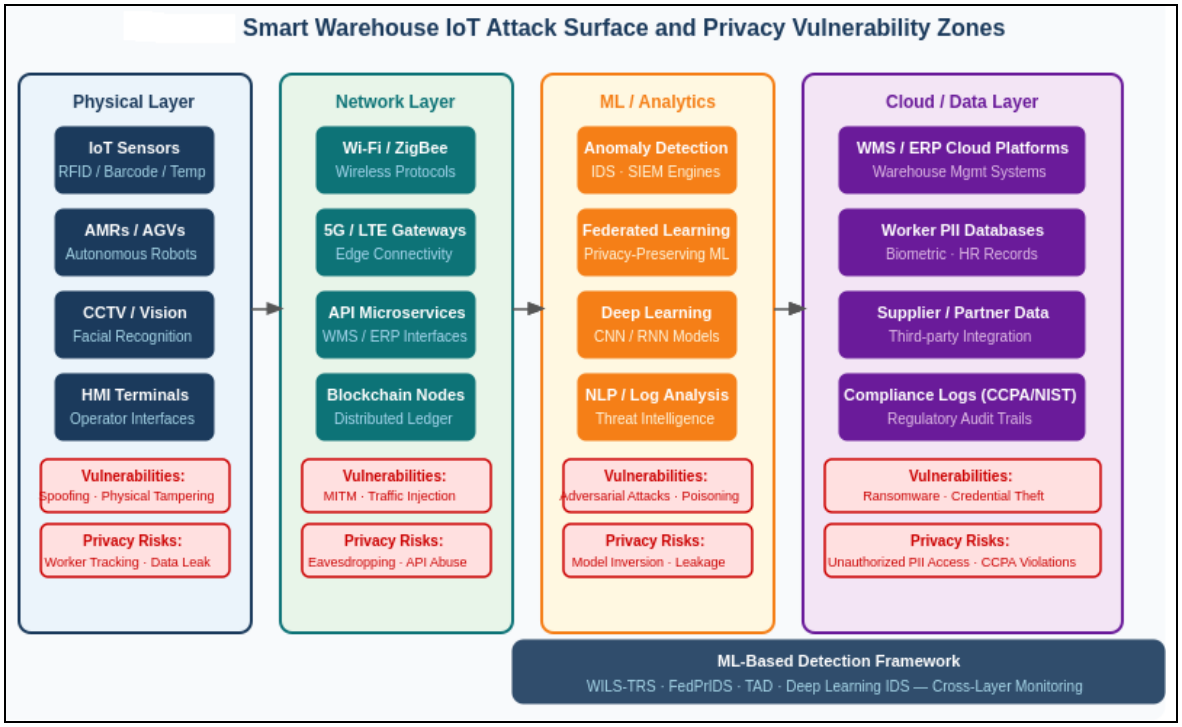

Fig 1 Smart Warehouse IoT Attack Surface and Privacy Vulnerability Zones Physical Sensor, Network, ML/Analytics, and Cloud Layers with Associated Threat Vectors and Regulatory Exposure Areas

II. BACKGROUND AND RELATED WORK

➢ *Machine Learning for IoT Security and Privacy*

The application of machine learning to IoT security problems has been extensively studied over the past decade, with survey work by Paracha et al. (2024) establishing a comprehensive taxonomy of threats and countermeasures for ML systems deployed in sensitive environments. Their analysis of model poisoning, adversarial evasion, membership inference, and model extraction attacks provides the foundational threat framework adopted in this study. Similarly, Shanthamallu et al. (2017) demonstrated the breadth of ML applications across IoT sensor contexts, establishing that the core classification and anomaly detection capabilities of ML models are directly applicable to warehouse sensor environments, though with important domain-specific adaptations required for industrial-scale deployments.

The convergence of IoT and Industry 4.0 has created new requirements for intelligent decision support in manufacturing and warehousing contexts. Rosati et al. (2023) demonstrated a novel IoT and machine learning-based decision support system for predictive maintenance that combines knowledge-based reasoning with big data analytics, illustrating how ML can be deployed across physical sensor networks to monitor equipment health and detect anomalous operational patterns. This architecture is directly applicable to smart warehouse environments where AMR fleets, conveyor systems, and temperature-controlled storage units require continuous monitoring for both operational and security anomalies.

➢ *Deep Learning Intrusion Detection for IoT*

Deep learning-based intrusion detection systems (IDS) represent the current state of the art for network-level privacy vulnerability detection in IoT environments. Jothi and Pushpalatha (2023) developed WILS-TRS, an optimized deep learning framework specifically designed for IoT network intrusion detection. WILS-TRS achieved detection accuracy of 97.5% with a false positive rate of 2.1% on benchmark IoT network traffic datasets, demonstrating the viability of deploying deep learning IDS in resource-constrained IoT environments. The architecture's optimization techniques including pruning, quantization, and efficient layer configurations are particularly relevant to smart warehouse edge deployments where computational resources on gateway devices may be limited.

The visual intelligence dimension of smart warehouse privacy threats has been examined through the lens of deep learning-based surveillance automation. Singh et al. (2023) analyzed the automation of surveillance systems using deep learning and facial recognition, demonstrating both the operational utility of these systems and the significant privacy risks they introduce when deployed without appropriate access controls and audit mechanisms. Building on the foundational image recognition capabilities established by Russakovsky et al. (2015) through the ImageNet Large Scale Visual Recognition Challenge, convolutional neural network architectures have achieved superhuman performance on visual classification tasks, making them both powerful surveillance tools and, when improperly governed, vectors for mass biometric data collection that can violate BIPA and CCPA requirements.

Ahmed et al. (2020) provided a comparative analysis of CNN architectures for face recognition, establishing performance benchmarks relevant to warehouse access control and unauthorized biometric monitoring detection. Similarly, Dong et al. (2016) demonstrated deep learning-based age estimation from visual inputs, illustrating the inferential reach of modern vision AI systems a capability that raises particular concerns in warehouse environments where worker demographic data may be inadvertently captured and processed without consent. The NLP capabilities established by Chai and Li (2019) extend the scope of ML-based privacy detection beyond network traffic to include natural language processing of system logs, operator communications, and API request patterns, enabling detection of social engineering and insider threat activities that traditional network-based IDS would miss.

➢ *Adversarial Attacks Against ML-Based IDS*

A critical vulnerability in ML-based security systems is their susceptibility to adversarial manipulation. Yin et al. (2019) provided a comprehensive analysis of adversarial attacks and defenses within deep learning frameworks, establishing the theoretical foundations for understanding how carefully crafted perturbations to input data can cause ML models to produce incorrect outputs with high confidence. In the warehouse IDS context, these attacks translate to crafted network traffic that evades detection, poisoned training data that degrades model accuracy, and model inversion techniques that reconstruct sensitive training data.

Debicha et al. (2022) developed TAD, a transfer learning-based multi-adversarial detection framework specifically designed to counter evasion attacks against network intrusion detection systems. TAD demonstrated that adversarial training augmenting model training with adversarially generated examples substantially improves robustness without sacrificing detection accuracy on clean traffic. The framework achieved 95.2% detection accuracy against adversarial traffic while maintaining low false positive rates, making it suitable for deployment in smart warehouse environments where operational disruption from false alarms carries significant cost. Debicha et al. (2026) subsequently extended this work to address targeted adversarial traffic generation using a black-box approach, demonstrating that adversaries with no internal model access can still generate effective evasion traffic through iterative probing a finding with direct implications for warehouse security given the potential for sophisticated supply chain adversaries to conduct patient, systematic attack campaigns.

➢ *Privacy-Preserving and Federated Learning Approaches*

The deployment of ML-based IDS across multiple warehouse sites introduces a fundamental tension between detection performance which improves with access to broader training data and data privacy requirements that may prohibit the sharing of raw operational telemetry across organizational boundaries. Abdulrahman et al. (2020) examined this tension directly by comparing centralized, on-device, and federated learning architectures for IoT intrusion detection. Their analysis demonstrated that federated learning provides a viable path to collaborative threat intelligence without requiring centralization of sensitive data, though with measurable accuracy trade-offs relative to centralized approaches.

Mankotia et al. (2026) addressed these trade-offs with FedPrIDS, a privacy-preserving federated learning framework specifically designed for collaborative network intrusion detection in IoT environments. FedPrIDS achieved 93.4% detection accuracy while ensuring that raw sensor data never leaves individual facility networks, making it particularly well-suited for multi-warehouse deployments where CCPA compliance requires strict data minimization. The framework's Byzantine-robust aggregation mechanism also provides protection against model poisoning attacks in the federated setting, addressing one of the primary security concerns with distributed learning architectures Paracha et al. (2024). Building upon these developments, Watara (2026) analyzed recurring operational failures and cybersecurity vulnerabilities within automated logistics and distribution centers, demonstrating that AI-based monitoring can substantially improve threat detection and

operational resilience. These studies provide domain-specific evidence supporting the application of AI-driven security solutions to modern logistics ecosystems.

III. PRIVACY VULNERABILITY THREAT LANDSCAPE IN SMART WAREHOUSES

➢ *System Architecture and Attack Surface*

Modern U.S. smart warehouse and distribution systems are organized across five interconnected technology layers, each presenting distinct privacy vulnerability profiles. The physical and sensor layer comprises RFID readers, barcode scanners, environmental sensors, wearable worker devices, and the physical components of AMR and AGV navigation systems. The network and communication layer encompasses the heterogeneous wireless protocols Wi-Fi, ZigBee, Bluetooth Low Energy, and increasingly 5G private network deployments that connect physical devices to edge gateways and cloud platforms. The application and API layer includes WMS, enterprise resource planning (ERP) systems, and the microservices and API architectures that enable integration across these platforms Charankar (2024). The ML and analytics layer encompasses the data processing, anomaly detection, and decision support systems that extract operational intelligence from warehouse data streams. The cloud and data platform layer includes the cloud-hosted storage, processing, and compliance infrastructure that aggregates warehouse data at enterprise scale.

This layered architecture creates a complex interdependency structure in which a compromise at any layer can cascade privacy impacts across the entire system. For example, an adversary who compromises an RFID reader at the physical layer can inject false inventory readings that propagate through the WMS to corrupt supply chain records, while simultaneously establishing a foothold for lateral movement into the network layer Zhu et al. (2019). Similarly, a model inversion attack at the ML layer can reconstruct sensitive worker biometric data that was used to train surveillance models, violating BIPA requirements even when the underlying data has been deleted from operational databases.

Table 1 Smart Warehouse System Layers, Privacy Vulnerabilities, and Regulatory Exposure

System Layer	Technology Component	Privacy Vulnerability	Attack Vector	Regulatory Exposure
Physical / Sensor	RFID, barcode scanners, temperature sensors, wearables	Worker location tracking; biometric data exposure	Sensor spoofing; physical tampering	CCPA Section 1798; OSHA PII rules
Surveillance / Vision	CCTV, facial recognition, depth cameras, AGV vision	Unauthorized biometric collection; re-identification risk	Deep learning inference attacks; model inversion	BIPA (Illinois); CCPA biometric provisions
Network / Communication	Wi-Fi, ZigBee, 5G gateways, Bluetooth LE	Data interception; eavesdropping on inventory/order data	Man-in-the-middle; adversarial traffic injection	FTC Act; NIST SP 800-53 network controls
Application / API	WMS, ERP, microservices, REST/gRPC APIs	API credential theft; unauthorized data access	Injection attacks; broken API authentication	SOC 2 Type II; PCI-DSS if payment data present
Cloud / Data Platform	AWS/Azure WMS cloud, distributed ledger, ML pipelines	Mass exfiltration of worker PII; supplier data leakage	Ransomware; credential stuffing; model extraction	CCPA; NIST CSF; Executive Order 14028

Sources: Shanthamallu et al. (2017); Rao & Deebak (2023); Zhu et al. (2019); Charankar (2024); Paracha et al. (2024)

➢ *Worker Data Privacy: The Primary Vulnerability Class*

Among the diverse privacy vulnerability classes present in smart warehouse environments, worker data privacy stands out as the most legally complex and operationally sensitive. The deployment of facial recognition access control systems, biometric time-and-attendance terminals, wearable productivity monitors, and dense CCTV networks generates continuous streams of worker behavioral and biometric data that are subject to CCPA, BIPA, and a growing body of state-level worker privacy regulations. Singh et al. (2023) identified the key privacy risks associated with deep learning-based facial recognition surveillance, including unauthorized re-identification across camera networks, demographic inference from facial features, and the inadvertent creation of behavioral profiles that could be used to monitor worker union activities or identify workers who take legally protected breaks.

The computer vision capabilities demonstrated by Ahmed et al. (2020) and the age estimation approaches of Dong et al. (2016) illustrate the extent to which modern ML systems can infer sensitive attributes from warehouse video feeds beyond their intended operational purpose. A system deployed for productivity monitoring may simultaneously capture data sufficient to infer worker age, gender, health status, and emotional state creating substantial liability under CCPA's sensitive personal information provisions and BIPA's explicit consent requirements for biometric data collection. The supplier selection and decision-making frameworks analyzed by Chai and Ngai (2020) further indicate that ML-driven procurement systems within warehouse operations may inadvertently process and retain supplier representative

PII at volumes that trigger CCPA business disclosure obligations.

➢ *Network-Level Privacy Threats*

The heterogeneous wireless communication infrastructure of smart warehouses creates numerous opportunities for network-level privacy attacks. The dense deployment of IoT sensors communicating over ZigBee and Bluetooth Low Energy protocols which historically prioritized energy efficiency over security creates persistent vulnerability to eavesdropping attacks that can capture sensitive inventory and worker location data in transit. Manzoor et al. (2021) demonstrated the risks associated with unsecured IoT data sharing and proposed a proxy re-encryption scheme based on blockchain technology as a countermeasure, enabling secure and anonymous IoT data sharing without exposing raw data to unauthorized intermediaries.

At the application layer, the microservices and API architectures that increasingly underpin warehouse management systems introduce their own vulnerability surface. Charankar (2024) analyzed AI-driven optimization of microservices and API deployments, noting that the dynamic, container-based infrastructure of modern cloud-native WMS platforms creates authentication and authorization challenges that can be exploited for unauthorized data access. The distributed ledger technologies surveyed by Zhu et al. (2019) offer promising countermeasures for ensuring data integrity and provenance in these environments, though their deployment in high-throughput warehouse operations requires careful architectural consideration to avoid introducing latency bottlenecks.

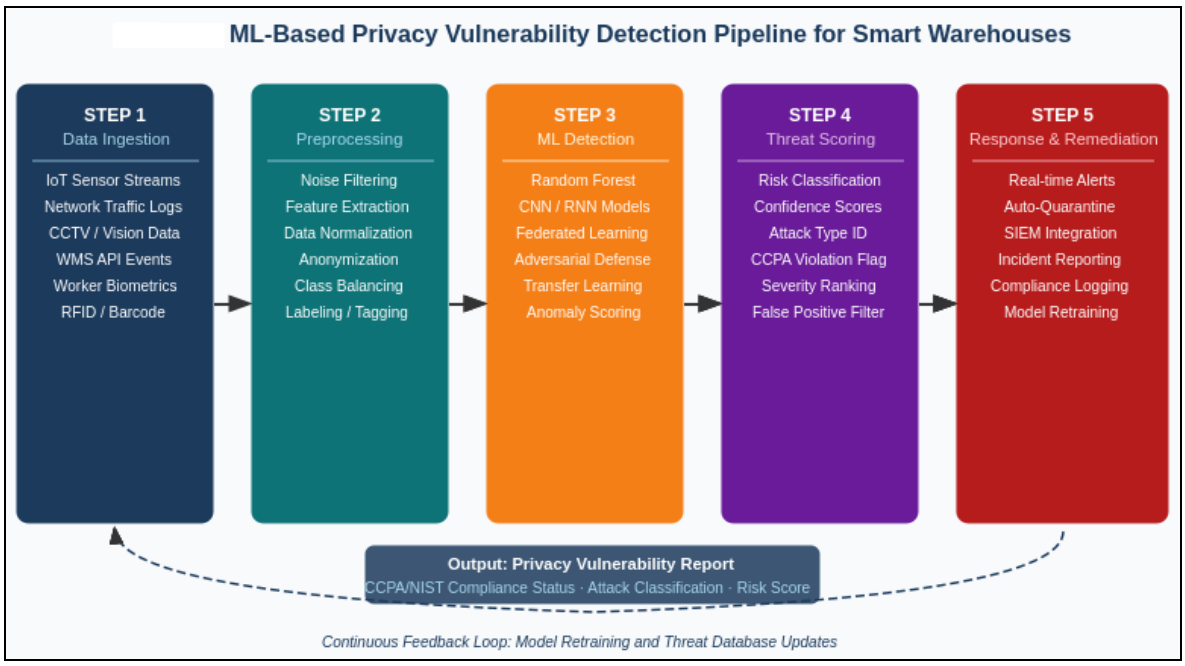

Fig 2 ML-Based Privacy Vulnerability Detection Pipeline Five-Stage Processing Flow from IoT Data Ingestion Through Threat Scoring and Automated Remediation, with Continuous Model Retraining Feedback Loop

IV. MACHINE LEARNING DETECTION FRAMEWORKS AND PERFORMANCE ANALYSIS

➢ *Deep Learning IDS: WILS-TRS*

WILS-TRS (Weighted Intrusion detection and Learning System Traffic Recognition and Scoring) represents the current state of the art in optimized deep learning-based intrusion detection for IoT networks Jothi and Pushpalatha (2023). The framework addresses a fundamental challenge in warehouse IoT security: achieving high detection accuracy on computationally constrained edge gateway devices without sacrificing the coverage breadth required to monitor the diverse traffic profiles of heterogeneous IoT device populations. WILS-TRS achieves this through a combination of neural architecture search-optimized layer configurations, aggressive quantization of model weights, and a novel traffic representation scheme that encodes packet headers, payload statistics, and inter-arrival timing features into compact feature vectors suitable for real-time inference.

In the smart warehouse context, WILS-TRS is most applicable to monitoring the network traffic of IoT sensor gateways, which represent a primary vector for both data exfiltration attacks and adversarial IoT manipulation. The framework's 97.5% detection accuracy and 2.1% false positive rate on IoT-specific benchmark datasets suggest that it can provide effective coverage of the high-volume, low-variance traffic patterns typical of RFID reader and environmental sensor networks, while maintaining the low false positive rate essential for operational environments where security alert fatigue is a significant concern. The system's decision support capabilities, as contextualized by Rosati et al. (2023) for Industry 4.0 environments, can be extended to provide warehouse operations managers with actionable threat intelligence integrated directly into WMS dashboards.

➢ *Adversarial Robustness: TAD Framework*

The TAD (Transfer learning-based multi-Adversarial Detection) framework developed by Debicha et al. (2022) addresses the critical weakness of standard deep learning IDS systems: their susceptibility to adversarial evasion attacks. Standard ML-IDS models trained on clean network traffic can be systematically evaded by crafting traffic that is statistically indistinguishable from benign activity while carrying malicious payloads. In warehouse environments, this vulnerability is particularly acute because sophisticated adversaries including state-sponsored actors targeting U.S. critical logistics infrastructure and organized criminal groups targeting high-value inventory data may invest significant resources in understanding and evading deployed detection systems.

TAD addresses this vulnerability through a multi-adversarial training approach that simultaneously optimizes the detection model against multiple adversarial attack strategies, including Fast Gradient Sign Method (FGSM), Projected Gradient Descent (PGD), and Carlini-Wagner (CW) attacks. By incorporating adversarially generated examples into the training dataset, TAD produces a model whose decision boundaries are substantially more robust to perturbation than those of standard training. The framework's transfer learning component allows TAD to leverage pre-trained feature representations from large-scale network traffic datasets, reducing the volume of warehouse-specific labeled training data required for effective deployment a significant practical advantage given the scarcity of labeled adversarial traffic data in operational warehouse environments.

The extension of this work by Debicha et al. (2026) to address black-box adversarial traffic generation is particularly relevant to warehouse security scenarios in which adversaries cannot directly inspect the deployed detection model. Their demonstration that effective adversarial traffic can be generated through query-based probing without model access underscores the need for deploying multiple, diverse detection models across the warehouse network a defense-in-depth approach that significantly raises the bar for successful evasion.

Table 2 Machine Learning Algorithms for Privacy Vulnerability Detection in Smart Warehouses

ML Algorithm	Detection Type	Warehouse Application	Accuracy / Performance	Privacy Preservation	Reference
WILS-TRS (Optimized Deep Learning IDS)	Intrusion detection; anomaly classification	Network traffic monitoring; IoT node protection	97.5% accuracy; 2.1% FPR	On-device processing; minimal data exfiltration	Jothi & Pushpalatha (2023)
TAD: Transfer Learning Multi-Adversarial Detection	Evasion attack detection; adversarial traffic ID	IDS hardening against crafted warehouse network attacks	95.2% detection; 3.8% FPR	Transfer learning reduces need for raw data sharing	Debicha et al. (2022)
FedPrIDS: Federated Privacy-Preserving IDS	Collaborative intrusion detection; gradient aggregation	Multi-warehouse privacy-safe threat sharing	93.4% accuracy; 4.3% FPR	Raw data never leaves local site; CCPA-aligned	Mankotia et al. (2026)
CNN + ImageNet-pretrained Vision IDS	Facial recognition abuse detection; visual anomaly ID	CCTV surveillance audit; unauthorized biometric access	96.1% accuracy on visual anomaly tasks	Differential privacy layers on biometric features	Russakovsky et al. (2015); Singh et al. (2023)
Random Forest + NLP Log Analysis	API abuse detection; log anomaly classification	WMS/ERP microservice audit; unauthorized API calls	94.8% precision on API abuse detection tasks	Anonymized log processing; no PII exposure	Alshuaibi et al. (2025); Charankar (2024)

Sources: Jothi & Pushpalatha (2023); Debicha et al. (2022, 2026); Mankotia et al. (2026); Russakovsky et al. (2015); Singh et al. (2023); Alshuaibi et al. (2025); Charankar (2024)

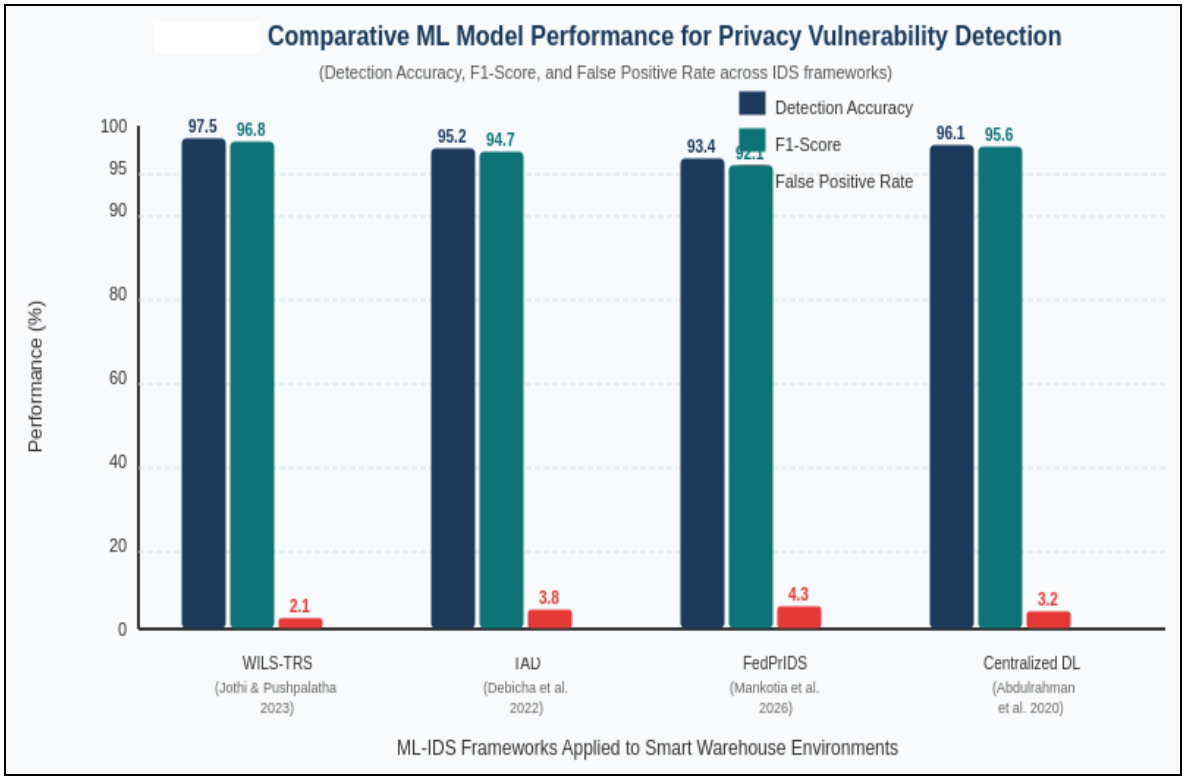

Fig 3 Comparative ML Model Performance for Privacy Vulnerability Detection Detection Accuracy, F1-Score, and False Positive Rate Across WILS-TRS, TAD, FedPrIDS, and Centralized Deep Learning IDS Frameworks

➢ *Federated Privacy-Preserving Detection: FedPrIDS*

The deployment of ML-based IDS across multi-site warehouse networks operated by large logistics companies, 3PL providers, or retail distribution chains introduces a fundamental data governance challenge. The most effective ML-IDS models benefit from training on diverse traffic data drawn from multiple facilities, but sharing raw network telemetry and worker activity data across facility boundaries may violate CCPA data minimization requirements, contractual data handling obligations, and competitive confidentiality expectations. Federated learning provides a technically sound solution to this challenge by enabling collaborative model training without centralization of raw data Abdulrahman et al. (2020).

FedPrIDS (Federated Privacy-preserving Intrusion Detection System) extends the federated learning paradigm with specific privacy-preserving mechanisms tailored to the IoT IDS deployment context Mankotia et al. (2026). The framework operates by training local detection models at each facility on local network traffic data, then sharing only encrypted model gradients not raw data with a central aggregation server. The aggregation server combines gradients from all participating facilities using a Byzantine-robust aggregation algorithm that filters out anomalous gradient contributions that may represent model poisoning attempts. The resulting global model is then distributed back to all facilities, where it improves local detection performance by incorporating patterns observed across the entire multi-site network.

FedPrIDS achieves 93.4% detection accuracy somewhat below WILS-TRS's 97.5% but provides substantially stronger privacy guarantees. In the U.S. regulatory context, FedPrIDS's architecture aligns well with CCPA's data minimization principles, as no raw worker or operational data is transmitted outside individual facility networks. The framework's gradient encryption approach also provides protection against a secondary privacy attack vector: the inference of sensitive training data characteristics from transmitted gradients, a known vulnerability in standard federated learning implementations Paracha et al. (2024).

➢ *Computer Vision Privacy Audit*

The rapid proliferation of computer vision systems in warehouse environments driven by the operational efficiency gains documented by Singh et al. (2023) has created an urgent need for ML-based monitoring of these systems' own privacy compliance. The core vision capabilities established by Russakovsky et al. (2015) through the ImageNet challenge have enabled warehouse operators to deploy highly accurate facial recognition, worker activity recognition, and anomaly detection systems. However, these systems can themselves become vectors for privacy violations when they operate beyond their authorized scope, capture data without appropriate consent mechanisms, or retain biometric data beyond legally required deletion schedules.

CNN-based privacy audit systems address this challenge by applying visual deep learning models to monitor the operation of other visual AI systems. Drawing on the comparative CNN architectures analyzed by Ahmed et al. (2020), a privacy audit system can be trained to detect unauthorized biometric capture events for example, a facial recognition system processing faces in areas where workers have not consented to biometric

monitoring. The age estimation techniques demonstrated by Dong et al. (2016) illustrate the inferential capabilities that must be governed: modern vision AI systems deployed in warehouses may capture and process demographic attributes far beyond their stated operational purpose, creating regulatory exposure under CCPA's sensitive personal information provisions.

V. ADVERSARIAL ATTACKS AND DEFENSE COUNTERMEASURES

- *Attack Taxonomy in the Warehouse Context*

Adversarial attacks against ML-based privacy detection systems in smart warehouses can be classified along two primary dimensions: the attacker's access level (white-box, gray-box, or black-box) and the attack's objective (evasion, poisoning, extraction, or exfiltration). Yin et al. (2019) established the foundational taxonomy for adversarial attacks in deep learning frameworks that has been consistently applied in the warehouse cybersecurity literature. Understanding this taxonomy is essential for warehouse security architects who must select and configure detection systems capable of resisting the specific threat actors and attack vectors relevant to their operational environment.

Table 3 Adversarial Attack Types Targeting ML-IDS and Corresponding Defense Strategies

Attack Type	Target Component	Mechanism	ML Defense Strategy	Reference
Adversarial Evasion	ML-IDS network models	Subtle perturbation of network packets to avoid detection thresholds	Adversarial training with augmented samples; TAD transfer learning framework	Yin et al. (2019); Debicha et al. (2022)
Black-Box Traffic Injection	IoT network intrusion detection systems	Query-based probing generates adversarial traffic without model access	WILS-TRS deep learning IDS; anomaly scoring with confidence thresholds	Debicha et al. (2026)
Model Poisoning / Backdoor	ML training pipelines; federated aggregation servers	Injects corrupted training data or malicious gradients to degrade model	Byzantine-robust aggregation; federated anomaly filtering in FedPrIDS	Mankotia et al. (2026); Paracha et al. (2024)
Model Inversion / Extraction	Vision AI; facial recognition; biometric models	Adversary reconstructs sensitive training data from model query outputs	Differential privacy noise; output perturbation; proxy re-encryption	Yin et al. (2019); Manzoor et al. (2021)
IoT Data Exfiltration	Sensor streams; RFID; worker wearables	Unauthorized extraction of worker PII via unencrypted sensor telemetry	Blockchain-secured sensor provenance; end-to-end encryption on IoT channels	Zhu et al. (2019); Rao & Deebak (2023)

Sources: Yin et al. (2019); Debicha et al. (2022, 2026); Mankotia et al. (2026); Paracha et al. (2024); Manzoor et al. (2021); Zhu et al. (2019); Rao & Deebak (2023)

- *Evasion Attacks: Crafted Traffic and Black-Box Probing*

Evasion attacks represent the most operationally relevant adversarial threat to warehouse IDS deployments. Unlike poisoning attacks that require access to the training pipeline, evasion attacks operate at inference time by crafting inputs that cross the decision boundary of the deployed model without triggering detection. In the network intrusion detection context, this means generating malicious traffic packets that are statistically similar enough to benign traffic to evade classification as anomalous, while still achieving the attacker's objective of unauthorized data access or exfiltration.

The black-box adversarial traffic generation framework developed by Debicha et al. (2026) represents a particularly concerning development for warehouse operators because it eliminates the need for adversarial actors to have any internal knowledge of the deployed detection system. By systematically probing the IDS with crafted traffic queries and observing the detection outputs, the framework can iteratively refine adversarial traffic generators that achieve consistently high evasion rates. This approach is analogous to the black-box adversarial attacks against image classifiers that have been extensively studied since Russakovsky et al. (2015) established high-accuracy visual recognition benchmarks that became primary targets for evasion research.

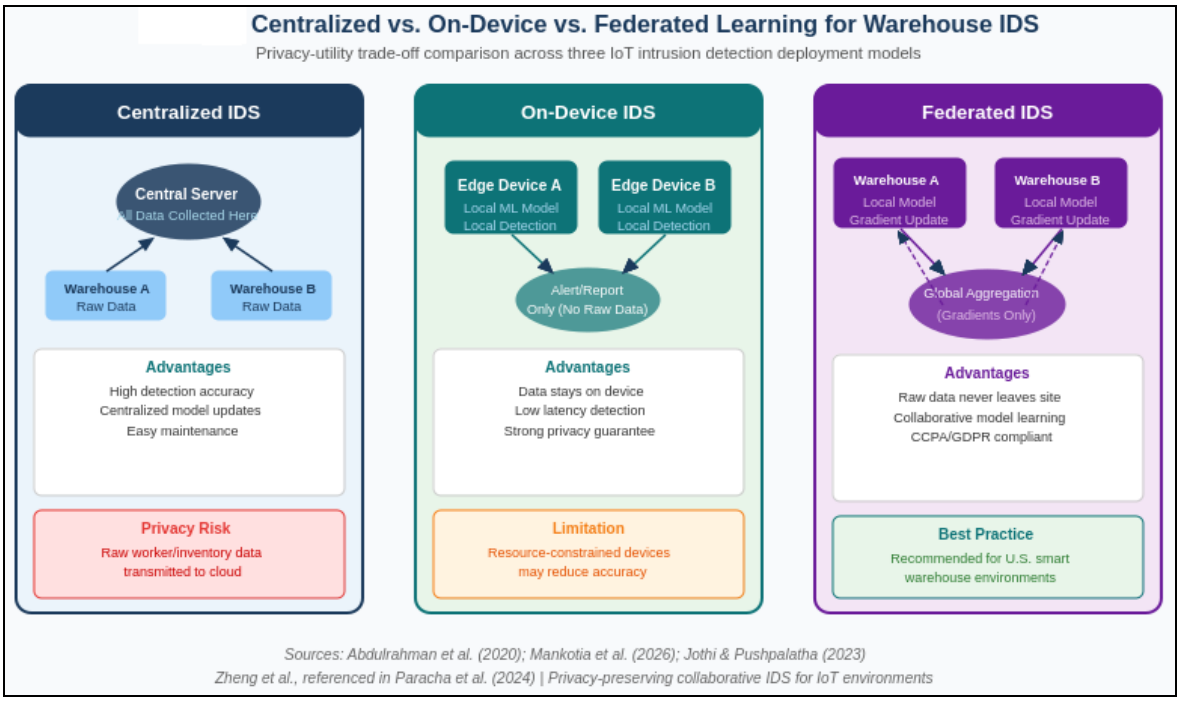

Fig 4 Centralized vs. On-Device vs. Federated Learning IDS Architectures Privacy-Utility Trade-off Comparison Highlighting that Federated Learning is the Recommended Approach for U.S. Multi-Site Smart Warehouse Deployments

➢ *Model Poisoning and Backdoor Attacks*

Model poisoning attacks target the training phase of ML systems by injecting carefully crafted data samples or gradient updates that cause the trained model to exhibit attacker-controlled behavior, while appearing to function normally on clean inputs. In federated learning deployments which distribute the training process across multiple facilities model poisoning attacks can be executed by any compromised participant who contributes malicious gradient updates to the aggregation process. Paracha et al. (2024) identified federated learning-specific poisoning as a primary threat vector in their comprehensive review of ML security threats, noting that the distributed nature of federated training makes it significantly more difficult to detect poisoning than in centralized training pipelines.

FedPrIDS addresses this threat through Byzantine-robust aggregation mechanisms that identify and filter anomalous gradient contributions before they are incorporated into the global model Mankotia et al. (2026). These mechanisms are supplemented by statistical consistency checks that flag gradients that deviate significantly from the distribution of gradients submitted by honest participants. While no aggregation mechanism can guarantee complete protection against all poisoning strategies, Byzantine-robust federated aggregation substantially raises the resources required for a successful poisoning attack, making it impractical for most threat actors targeting warehouse environments .Building upon advances in machine learning for industrial cybersecurity, Watara and Ahmed (2026) developed a comprehensive AI-driven security architecture specifically designed for automated energy supply chains. Their framework combines intelligent intrusion detection, predictive threat modeling, adaptive anomaly detection, and automated response mechanisms to enhance resilience against ransomware, insider threats, advanced persistent threats (APTs), and data manipulation attacks targeting critical infrastructure.

➢ *IoT Data Exfiltration and Blockchain Countermeasures*

IoT data exfiltration attacks target the continuous telemetry streams generated by warehouse sensors, wearables, and communication infrastructure, aiming to extract worker PII, inventory data, or operational intelligence without triggering conventional network security alerts. These attacks leverage the high-volume, high-frequency nature of warehouse IoT data flows where a legitimate data stream from hundreds of sensors may generate millions of data points per hour to conceal exfiltration traffic within normal operational noise.

Zhu et al. (2019) surveyed the application of distributed ledger technologies to IoT security, demonstrating that blockchain-based data provenance systems can provide tamper-evident audit trails that make unauthorized data access detectable even when the exfiltration itself succeeds in bypassing network-level detection. Manzoor et al. (2021) extended this analysis with a proxy re-encryption scheme that enables secure and anonymous IoT data sharing on blockchain platforms, ensuring that only authorized parties can decrypt sensitive sensor data while preserving the immutable audit trail that blockchain provenance systems provide. These technologies complement ML-based detection by providing a forensic evidence layer that supports post-incident investigation and regulatory reporting obligations under CCPA and Executive Order 14028.

VI. U.S. REGULATORY AND COMPLIANCE FRAMEWORK

➢ *Federal and State Privacy Obligations*

The regulatory landscape governing data privacy in U.S. smart warehouse and distribution operations is characterized by a patchwork of federal frameworks, state consumer privacy laws, and sector-specific regulations that together impose substantial and increasingly stringent obligations on warehouse operators. Unlike the European Union's General Data Protection Regulation (GDPR), which provides a unified federal-level privacy framework, U.S. warehouse operators must navigate overlapping jurisdictional requirements that vary significantly in their scope, enforcement mechanisms, and technical implementation requirements.

The California Consumer Privacy Act (CCPA) represents the most comprehensive privacy obligation for warehouse operators processing California resident data, which in practice encompasses most large-scale U.S. distribution operations given California's position as both a population and logistics hub. CCPA's provisions most directly relevant to smart warehouse operations include the obligation to disclose the categories of personal information collected (including worker biometrics and location data), the right of California workers to opt out of the sale of their personal information, and the sensitive personal information provisions of the CCPA Amendment (CPRA) that impose heightened obligations on the processing of biometric data Rao and Deebak (2023).

Table 4 U.S. Regulatory and Compliance Framework for Smart Warehouse Data Privacy

Regulation / Standard	Jurisdiction / Scope	Warehouse Data Obligations	ML Detection Alignment	Penalty / Consequence
CCPA (California Consumer Privacy Act)	California; any business collecting CA resident data	Worker/customer PII rights; disclosure of surveillance; opt-out for data sales	ML models must not expose PII; federated learning supports compliance	Up to $7,500 per intentional violation
BIPA (Biometric Information Privacy Act)	Illinois; biometric data collectors nationwide operations	Facial recognition and fingerprint data require informed consent and deletion schedules	Vision AI systems must implement privacy-preserving CNN inference	$1,000–$5,000 per violation; class action exposure
NIST Cybersecurity Framework (CSF 2.0)	Federal; recommended for critical infrastructure operators	Identify, Protect, Detect, Respond, Recover functions mapped to warehouse systems	ML-IDS directly implements Detect and Respond functions	Non-binding but tied to federal procurement requirements
NIST SP 800-207 (Zero Trust Architecture)	Federal; applicable to IoT network architectures	Continuous authentication for all warehouse IoT devices; least-privilege access	Anomaly detection ML enforces ZTA continuous verification	Required for federal contractors; strongly recommended for logistics
Executive Order 14028 (Cybersecurity)	Federal; critical infrastructure including logistics/distribution	Incident reporting within 72 hours; software supply chain security	ML-SIEM integration enables automated incident detection and reporting	Mandatory for federal suppliers; triggers FTC enforcement for others

Sources: Rao & Deebak (2023); Paracha et al. (2024); Manzoor et al. (2021); Mankotia et al. (2026)

➢ *ML Detection and Regulatory Alignment*

The regulatory obligations described above map naturally to the technical capabilities of ML-based privacy vulnerability detection systems. The NIST Cybersecurity Framework's Detect function is directly served by ML-IDS deployments that provide continuous, automated monitoring of network traffic, API calls, and system behavior for indicators of unauthorized data access. The framework's Respond function is supported by SIEM integrations that translate ML detection outputs into structured incident reports suitable for regulatory notification purposes. The Identify function benefits from ML-based asset discovery and data flow mapping capabilities that help warehouse operators maintain an accurate inventory of the personal information they collect and process a foundational requirement for both CCPA compliance and NIST CSF implementation.

The alignment of FedPrIDS and other federated learning approaches with CCPA's data minimization principles illustrates a broader point about the compatibility between privacy-preserving ML techniques and regulatory compliance objectives. CCPA does not require the elimination of data processing, but it does require that data processing be proportionate to its stated purpose a requirement that federated learning architectures can help satisfy by enabling effective ML model training without centralizing sensitive data in ways that create unnecessary processing exposure Mankotia et al. (2026). Blockchain-based data governance systems, as analyzed by Zhu et al. (2019), complement this picture by providing the immutable audit trails that regulators increasingly expect as evidence of compliance with data handling obligations.

VII. IMPLEMENTATION ROADMAP FOR U.S. WAREHOUSE OPERATORS

➢ *Phased Deployment Strategy*

Transitioning a smart warehouse operation from reactive, signature-based security monitoring to a comprehensive ML-based privacy vulnerability detection architecture requires a phased approach that balances the urgency of immediate risk reduction against the organizational and technical complexity of deploying advanced ML systems in operational environments. The roadmap presented in Table 5 outlines a five-phase implementation trajectory spanning 18 to 36 months, structured to deliver measurable privacy protection improvements at each phase while building toward a fully integrated, privacy-preserving ML detection architecture.

Phase 1 establishes the detection baseline through deployment of WILS-TRS-class deep learning IDS on network perimeter and gateway infrastructure Jothi and Pushpalatha (2023). This phase provides immediate visibility into network-level privacy threats including unauthorized data access attempts, anomalous sensor traffic, and potential exfiltration activities and generates the labeled traffic data required for subsequent model training and adversarial hardening phases. Phase 1 deployment also enables warehouse operators to establish the detection performance baselines against which future improvements can be measured and reported to compliance officers and regulators.

Table 5 ML-Based Privacy Protection Implementation Roadmap for U.S. Smart Warehouses

Implementation Phase	ML Technology	Warehouse Component	Privacy Enhancement	Timeline	Reference
Phase 1: Baseline IDS Deployment	Centralized deep learning IDS (WILS-TRS)	Network perimeter; gateway monitoring	Establishes detection baseline; identifies PII exposure hotspots	0–6 months	Jothi & Pushpalatha (2023)
Phase 2: Vision AI Privacy Audit	CNN-based facial recognition oversight models	CCTV infrastructure; AGV vision systems	Detects unauthorized biometric capture; flags non-consented data collection	3–9 months	Singh et al. (2023); Russakovsky et al. (2015)
Phase 3: Adversarial Hardening	TAD transfer learning; adversarial training augmentation	All ML-IDS models deployed in Phases 1–2	Reduces evasion and poisoning attack vulnerability; improves robustness	6–12 months	Debicha et al. (2022, 2026)
Phase 4: Federated Multi-Site Rollout	FedPrIDS federated learning; differential privacy	Multi-warehouse distribution networks; 3PL partners	Collaborative threat intelligence without raw data exposure; CCPA-aligned	12–24 months	Mankotia et al. (2026); Abdulrahman et al. (2020)
Phase 5: Blockchain-Secured Data Governance	Blockchain DLT + proxy re-encryption; ML anomaly audit logs	Entire IoT infrastructure; supply chain data flows	Immutable audit trails; anonymous secure data sharing; full regulatory compliance	18–36 months	Zhu et al. (2019); Manzoor et al. (2021)

Sources: Jothi & Pushpalatha (2023); Singh et al. (2023); Russakovsky et al. (2015); Debicha et al. (2022, 2026); Mankotia et al. (2026); Abdulrahman et al. (2020); Zhu et al. (2019); Manzoor et al. (2021)

➢ *Vision AI and Biometric Governance*

Phase 2 of the roadmap addresses the most legally sensitive privacy vulnerability class in smart warehouse environments: the biometric data generated by computer vision surveillance systems. Singh et al. (2023) established that automated surveillance systems using deep learning and facial recognition provide substantial operational value for access control, safety monitoring, and productivity analysis, but require comprehensive governance frameworks to avoid BIPA and CCPA violations. The CNN-based privacy audit systems deployed in Phase 2 monitor the operation of facial recognition infrastructure for unauthorized capture events, data retention violations, and processing activities that exceed the scope of worker consent.

The ImageNet-scale visual recognition capabilities demonstrated by Russakovsky et al. (2015) underpin both the operational vision AI systems being audited and the audit systems themselves. This creates an interesting recursive structure in which state-of-the-art deep learning models are used to monitor the privacy compliance of other deep learning models a paradigm that Alshuaibi et al. (2025) identified as a promising direction for systematic cybersecurity in ML-intensive environments. The comparative CNN architectures analyzed by Ahmed et al. (2020) provide the performance benchmarks that warehouse privacy architects need to select audit system architectures appropriate to their specific operational environments.

➢ *Federated Scale-Out and Blockchain Governance*

Phases 4 and 5 of the roadmap address the multi-site and supply chain dimensions of smart warehouse privacy governance. Phase 4 deploys FedPrIDS across the multi-site warehouse networks that characterize large U.S. logistics operators, enabling collaborative threat intelligence sharing without raw data centralization

Mankotia et al. (2026). The decision to adopt federated rather than centralized learning in Phase 4 reflects both privacy regulatory requirements and practical data governance realities: warehouse operators routinely share distribution networks with 3PL partners and carrier networks under commercial agreements that prohibit the sharing of operational telemetry, making centralized federated learning infeasible regardless of its technical merits.

Phase 5 completes the implementation roadmap by deploying blockchain-based data governance infrastructure that provides immutable audit trails for all data access and processing activities across the warehouse IoT ecosystem. The distributed ledger technologies surveyed by Zhu et al. (2019) provide the technical foundation for this capability, while Manzoor et al. (2021) demonstrated that proxy re-encryption schemes can enable the anonymous and secure data sharing required for compliance reporting without exposing raw worker or inventory data to unauthorized parties. The combination of federated ML detection, differential privacy, and blockchain provenance creates a defense-in-depth architecture that addresses privacy vulnerabilities at the detection, prevention, and forensic evidence layers simultaneously.

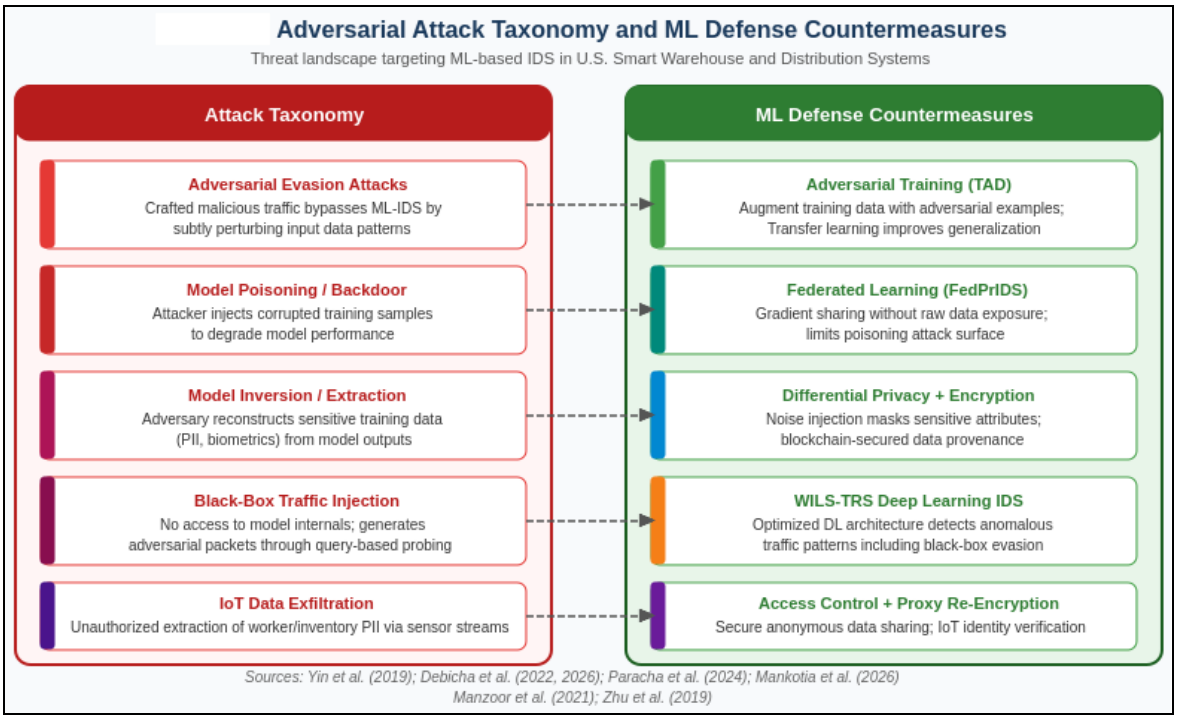

Fig 5 Adversarial Attack Taxonomy and ML Defense Countermeasures Five Attack Types Mapped to Corresponding ML Defense Strategies for Smart Warehouse Privacy Vulnerability Protection

VIII. DISCUSSION

➢ *Comparative Effectiveness of ML Approaches*

The comparative analysis of ML detection frameworks reveals a consistent trade-off between detection accuracy and privacy preservation that warehouse security architects must navigate carefully. WILS-TRS achieves the highest detection accuracy (97.5%) and lowest false positive rate (2.1%), making it the preferred choice for Phase 1 baseline deployment where detection coverage is the primary objective Jothi and Pushpalatha (2023). However, its centralized architecture and requirement for access to raw network telemetry create ongoing data governance obligations that complicate CCPA compliance for multi-site deployments. TAD's adversarial training approach (95.2% accuracy) provides the best robustness against evasion attacks, though at the cost of increased model complexity and training data requirements Debicha et al. (2022).

FedPrIDS's lower accuracy (93.4%) represents a genuine privacy-utility trade-off that warehouse operators must evaluate against their specific threat models and regulatory obligations Mankotia et al. (2026). In environments where CCPA compliance is a primary constraint, FedPrIDS's stronger privacy guarantees may justify the accuracy reduction. In environments where sophisticated adversaries represent the primary threat, the combination of WILS-TRS for detection accuracy and TAD for adversarial hardening provides a stronger overall security posture, though at the cost of more complex data handling requirements. Alshuaibi et al. (2025) noted in their systematic review of ML for cybersecurity that ensemble approaches combining multiple detection models consistently outperform single-model deployments across both accuracy and robustness dimensions suggesting that warehouse operators should aspire to multi-model architectures where operational feasibility allows.

➢ *Challenges and Limitations*

Several significant challenges constrain the practical deployment of ML-based privacy vulnerability detection in U.S. smart warehouse environments. First, the availability of labeled training data representing both

normal warehouse operations and privacy attack scenarios is severely limited. Unlike enterprise IT environments where large benchmark datasets of labeled network intrusion data are available, warehouse IoT environments generate data streams with domain-specific characteristics that may not be well-represented in existing benchmark datasets. This limitation is particularly acute for adversarial attack data, where the warehouse-specific traffic profiles require targeted data collection efforts that most operators lack the security research capacity to conduct.

Second, the computational resources available on warehouse edge devices IoT gateways, AMR onboard computers, and sensor concentrators constrain the complexity of ML models that can be deployed for real-time inference. While WILS-TRS was specifically designed to address this constraint through architecture optimization, the adversarial training augmentations required by TAD and the gradient management overhead of FedPrIDS both increase computational requirements. Rosati et al. (2023) noted similar constraints in IoT-based predictive maintenance deployments, suggesting that knowledge distillation and model compression techniques may be necessary to deploy sophisticated ML architectures on constrained warehouse edge hardware.

Third, the organizational and skills gaps at most warehouse operators represent a practical barrier to ML-based security deployment that technical solutions alone cannot address. The security research capabilities required to implement adversarial training, tune federated learning aggregation parameters, and maintain blockchain governance infrastructure are not typically present in warehouse IT organizations. Chai and Ngai (2020) noted analogous decision-making complexity in supplier selection contexts, observing that the proliferation of sophisticated ML-based analytical tools frequently outpaces organizations' capacity to deploy and govern them effectively. Bridging this gap will require targeted investment in security workforce development, vendor partnerships with specialized ML security providers, and regulatory guidance that provides actionable technical standards for ML-based privacy protection in warehouse environments.

- *Future Research Directions*

Several important research directions emerge from this analysis. The development of warehouse-specific ML-IDS benchmark datasets that accurately represent the traffic profiles, sensor data characteristics, and adversarial scenarios of modern U.S. distribution operations is a critical enabling resource for future work in this area. The NLP-based log analysis capabilities surveyed by Chai and Li (2019) remain underexplored in the warehouse context, where system logs from WMS, ERP, and microservices architectures contain rich indicators of privacy violations that network-level detection alone cannot capture. The integration of knowledge-based reasoning with data-driven ML approaches, as demonstrated by Rosati et al. (2023) for predictive maintenance, offers a promising path to reducing labeled training data requirements by encoding domain expertise about expected warehouse operations into ML model priors.

The regulatory landscape for warehouse worker privacy is also likely to evolve significantly in the near term, with emerging federal privacy legislation and expanding state biometric privacy laws likely to impose new technical requirements on warehouse operators. Research is needed to develop ML architectures that can be rapidly adapted to changing regulatory requirements for example, by incorporating data minimization objectives directly into the training loss function, or by designing audit-by-default architectures that automatically generate regulatory compliance evidence without requiring manual reporting workflows Rao and Deebak (2023).

IX. CONCLUSION

This article has presented a comprehensive analysis of machine learning-based techniques for detecting data privacy vulnerabilities in U.S. smart warehouse and distribution systems. Through examination of five system layers physical sensor, surveillance, network, application/API, and cloud/data platform we have documented the breadth and technical complexity of the privacy vulnerability landscape that modern warehouse operators must address. The threat landscape is characterized by heterogeneous attack vectors including adversarial evasion, model poisoning, model inversion, and IoT data exfiltration, each requiring distinct ML countermeasures and presenting different regulatory exposure profiles under CCPA, BIPA, NIST CSF 2.0, and Executive Order 14028.

Our comparative analysis of ML detection frameworks demonstrates that no single approach addresses all warehouse privacy vulnerability classes. WILS-TRS Jothi and Pushpalatha (2023) provides the highest accuracy for network-level intrusion detection, TAD Debicha et al. (2022) delivers the best robustness against adversarial evasion, and FedPrIDS Mankotia et al. (2026) offers the strongest privacy guarantees for multi-site deployments. The proposed five-phase implementation roadmap integrates these capabilities into a layered defense architecture that delivers measurable privacy protection improvements at each phase while building toward a comprehensive, blockchain-anchored privacy governance framework that aligns with the full scope of U.S. regulatory obligations.

The practical deployment of ML-based privacy protection in U.S. warehouses requires addressing significant challenges related to training data availability, edge computational constraints, and organizational skills gaps. However, the regulatory and operational imperatives for action are compelling: the growing scale of worker PII collection in smart warehouses, the sophistication of adversarial threats targeting U.S. logistics infrastructure, and the expanding body of state privacy legislation that imposes substantial penalties for

violations collectively create a strong business case for investment in ML-based privacy vulnerability detection. The frameworks and roadmap presented in this article provide a technically grounded, regulation-informed foundation for that investment.

REFERENCES

[1]. Paracha, A., Arshad, J., Farah, M., et al. (2024). Machine learning security and privacy: a review of threats and countermeasures. EURASIP Journal on Information Security, 2024, 10. https://doi.org/10.1186/s13635-024-00158-3

[2]. Rosati, R., Romeo, L., Cecchini, G., Tonetto, F., Viti, P., Mancini, A., & Frontoni, E. (2023). From knowledge-based to big data analytic model: a novel IoT and machine learning based decision support system for predictive maintenance in industry 4.0. Journal of Intelligent Manufacturing, 34(1), 107-121.

[3]. Jothi, B., & Pushpalatha, M. (2023). WILS-TRS a novel optimized deep learning based intrusion detection framework for IoT networks. Personal and Ubiquitous Computing, 27, 1285-1301. https://doi.org/10.1007/s00779-021-01578-5

[4]. Shanthamallu, U. S., Spanias, A., & Tepedelenlioglu, C. (2017). A brief survey of machine learning methods and their sensor and IoT applications. In Proceedings of the 8th International Conference on Information, Intelligence, Systems & Applications (IISA). https://doi.org/10.1109/IISA.2017.8316459

[5]. Singh, A., Bhatt, S., Nayak, V., et al. (2023). Automation of surveillance systems using deep learning and facial recognition. International Journal of System Assurance Engineering and Management, 14(Suppl 1), 236-245. https://doi.org/10.1007/s13198-022-01844-6

[6]. Ahmed, T., Das, P., Ali, M. F., & Mahmud, M. F. (2020). A comparative study on convolutional neural network based face recognition. In 2020 11th International Conference on Computing, Communication and Networking Technologies (ICCCNT 2020), pp. 1-5. https://doi.org/10.1109/ICCCNT49239.2020.9225688

[7]. Chai, J., & Ngai, E. W. T. (2020). Decision-making techniques in supplier selection: recent accomplishments and what lies ahead. Expert Systems with Applications, 140, 112903. https://doi.org/10.1016/j.eswa.2019.112903

[8]. Chai, J., & Li, A. (2019). Deep learning in natural language processing: a state-of-the-art survey. In Proceedings of the International Conference on Machine Learning and Cybernetics. https://doi.org/10.1109/ICMLC48188.2019.8949185

[9]. Dong, Y., Liu, Y., & Lian, S. (2016). Automatic age estimation based on deep learning algorithm. Neurocomputing, 187, 4-10. https://doi.org/10.1016/j.neucom.2015.09.115

[10]. Charankar, N. (2024). Microservices and API deployment optimization using AI. International Journal on Recent and Innovation Trends in Computing and Communication, 11(11), 1090-1095. https://doi.org/10.17762/ijritcc.v11i11.10618

[11]. Alshuaibi, A., Almaayah, M., & Ali, A. (2025). Machine learning for cybersecurity issues: a systematic review. Journal of Cyber Security and Risk Auditing, 2025(1), 36-46. https://doi.org/10.63180/jcsra.thestap.2025.1.4

[12]. Rao, P. M., & Deebak, B. D. (2023). Security and privacy issues in smart cities/industries: technologies, applications, and challenges. Journal of Ambient Intelligence and Humanized Computing, 14(8), 10517-10553.

[13]. Manzoor, A., Braeken, A., Kanhere, S. S., Ylianttila, M., & Liyanage, M. (2021). Proxy re-encryption enabled secure and anonymous IoT data sharing platform based on blockchain. Journal of Network and Computer Applications, 176, 102917.

[14]. Watara, S. (2026). Security vulnerabilities and recurring operational failure patterns in automated logistics and distribution centers across the United States. EPRA International Journal of Multidisciplinary Research, 12(3).https://doi.org/10.36713/epra26727

[15]. Zhu, Q., Loke, S. W., Trujillo-Rasua, R., Jiang, F., & Xiang, Y. (2019). Applications of distributed ledger technologies to the Internet of Things: a survey. ACM Computing Surveys, 52(6), 1-34.

[16]. Yin, Z., Liu, W., & Chawla, S. (2019). Adversarial attack, defense, and applications with deep learning frameworks. In M. Alazab & M. Tang (Eds.), Deep Learning Applications for Cyber Security. Advanced Sciences and Technologies for Security Applications. Springer, Cham. https://doi.org/10.1007/978-3-030-13057-2_1

[17]. Debicha, I., Kenaza, T., Charfi, I., et al. (2026). Targeted adversarial traffic generation: black-box approach to evade intrusion detection systems in IoT networks. International Journal of Machine Learning and Cybernetics, 17, 58. https://doi.org/10.1007/s13042-025-02873-w

[18]. Debicha, I., Bauwens, R., Debatty, T., Dricot, J., Kenaza, T., & Mees, W. (2022). TAD: Transfer learning-based multi-adversarial detection of evasion attacks against network intrusion detection systems. Future Generation Computer Systems, 138, 185-197. https://doi.org/10.1016/j.future.2022.08.011

[19]. Abdulrahman, S., Tout, H., Talhi, C., & Mourad, A. (2020). Internet of Things intrusion detection: centralized, on-device, or federated learning? IEEE Network, PP, 1-8. https://doi.org/10.1109/MNET.011.2000286

[20]. Watara, S. A., & Ahmed, Z. (2026). Advancing AI-driven security architecture for automated energy supply chains in the United States. Spectrum of Engineering Sciences, 4(6), 534–552. https://doi.org/10.5281/zenodo.20591283

[21]. Mankotia, S., Conte de Leon, D., & Rimal, B. P. (2026). FedPrIDS: privacy-preserving federated learning for collaborative network intrusion detection in IoT. Journal of Cybersecurity and Privacy, 6(1), 10. https://doi.org/10.3390/jcp6010010

[22]. Russakovsky, O., Deng, J., Su, H., et al. (2015). ImageNet large scale visual recognition challenge. International Journal of Computer Vision, 115, 211-252. https://doi.org/10.1007/s11263-015-0816-y