MACHINE LEARNING FOR NETWORK INTRUSION DETECTION IN USA CRITICAL INFRASTRUCTURE: CHALLENGES AND OPPORTUNITIES

Joy Selasi Agbesi[1], Abigail Nanayaa Otchill[2], Raymond Horlalie Tay[3] and Noah K. Bamfo[4]

[1]Department; J. Warren McClure School of Emerging Communication & Technology, Ohio University, USA
[2]Network Engineer, Foundation and Support, Meta, Richmond VA, United States
[3]College of Engineering Northeastern University, Boston, MA, United States
[4]Consulting Network Engineer, IT Department, Lidl US, Arlington VA, United States

ABSTRACT

The convergence of information technology and operational technology in United States critical infrastructure has created unprecedented efficiency gains while simultaneously expanding attack surfaces vulnerable to sophisticated cyber threats. This paper examines the application of machine learning to network intrusion detection in critical infrastructure, with particular emphasis on smart cities and power grid implementations. Through comprehensive analysis of current threat landscapes, technical approaches, and operational constraints, the study identifies key challenges impeding the deployment of machine learning-based security solutions, including data scarcity, class imbalance, concept drift, and adversarial robustness concerns.

The analysis reveals that while machine learning offers promising capabilities for detecting anomalous patterns and previously unknown attack vectors beyond traditional signature-based systems, successful implementation requires addressing fundamental tensions between real-time operational requirements and computational complexity, between model explainability and detection accuracy, and between privacy preservation and effective security monitoring. The paper examines specific vulnerabilities in smart grid architectures, municipal systems, and IoT-enabled infrastructure, demonstrating how heterogeneous device ecosystems and legacy system integration compound security challenges.

Furthermore, the study synthesizes emerging opportunities including ensemble detection approaches, physics-informed machine learning, transfer learning techniques, federated learning, explainable artificial intelligence, and collaborative threat intelligence sharing mechanisms. It proposes a framework for cross-sector collaboration and outlines standardized evaluation methodologies essential for validating machine learning security solutions in safety-critical environments. The findings indicate that realizing the full potential of machine learning for infrastructure protection requires coordinated efforts spanning technology development, workforce capacity building, regulatory framework evolution, and sustained information sharing across stakeholder communities. This work contributes to the growing body of knowledge on securing increasingly interconnected critical infrastructure systems upon which modern society fundamentally depends.

KEYWORDS

Machine learning, intrusion detection, smart cities, smart grid, IoT security, anomaly detection, operational technology, cybersecurity, federated learning, explainable AI

1. INTRODUCTION

The landscape of critical infrastructure protection in the United States has undergone a fundamental transformation as networks have evolved from isolated operational technology environments to highly interconnected systems spanning multiple domains. This convergence of information technology and operational technology, while enabling unprecedented efficiency and coordination, has simultaneously expanded the attack surface available to malicious actors. The emergence of smart cities represents a particularly compelling example of this transformation, where municipal services ranging from water treatment facilities to traffic management systems increasingly rely on complex networks of sensors, actuators, and control systems. As cities across America invest heavily in these interconnected infrastructures, the challenge of securing them against sophisticated cyber threats has become paramount. Within this context, machine learning has emerged as a promising approach for detecting network intrusions, offering the potential to identify anomalous patterns and previously unknown attack vectors that traditional signature-based systems might miss. However, the application of machine learning to critical infrastructure protection is fraught with unique challenges that stem from the operational constraints, safety requirements, and adversarial sophistication characteristic of these environments.

The intersection of critical infrastructure security and machine learning represents more than a technical challenge; it embodies a fundamental question about how societies can harness advanced computational techniques to protect the systems upon which modern life depends. Smart cities epitomize this challenge, as they integrate diverse infrastructures including power grids, transportation networks, water systems, and emergency services into cohesive digital ecosystems. Research has demonstrated that these interconnected systems face mounting security challenges as their complexity increases. The deployment of Internet of Things devices throughout municipal infrastructure has created millions of potential entry points for attackers, each representing a vulnerability that must be monitored and protected [1]. The scale of this challenge is staggering; modern smart city implementations can involve hundreds of thousands of connected devices generating massive volumes of network traffic that must be analyzed in real-time to detect malicious activity. Traditional security approaches based on predefined signatures and rule-based detection have proven inadequate for this environment, as they cannot adapt to the rapidly evolving tactics employed by sophisticated adversaries. Machine learning offers a potential solution by enabling systems to learn normal behavior patterns and identify deviations that may indicate intrusions, but realizing this potential requires overcoming significant technical and operational obstacles.

1.1. Contributions

The main contributions of this paper are summarized as follows:

1. **Critical infrastructure threat analysis:**
   Presents a structured examination of cyber threats affecting U.S. critical infrastructure, emphasizing IT–OT convergence, protocol vulnerabilities, and limitations of traditional intrusion detection systems.
2. **Machine learning–based intrusion detection taxonomy:**
   Provides a concise classification of supervised, unsupervised, deep learning, and hybrid approaches, analyzed under infrastructure-specific constraints including data imbalance, concept drift, and real-time operational requirements.
3. **Infrastructure-aware challenge assessment:**
   Identifies key technical and operational challenges such as adversarial robustness, model

explainability, privacy preservation, and safety-critical false-positive risks in smart grid and smart city environments.

4. **Future-oriented framework and research directions:**
   Synthesizes emerging solutions—including federated learning, explainable AI, ensemble methods, and physics-informed models—and outlines collaborative and evaluation considerations for trustworthy deployment.

2. THE CRITICAL INFRASTRUCTURE THREAT LANDSCAPE

Understanding the threat landscape facing American critical infrastructure requires examining both the evolving nature of cyber-attacks and the specific vulnerabilities inherent in industrial control systems and municipal networks. Recent vulnerability disclosures have highlighted the severity of risks facing networked infrastructure components. The discovery of CVE-2023-20080, affecting Cisco IOS and IOS XE Software IPv6 DHCP implementations, demonstrated how fundamental networking protocols can harbor critical vulnerabilities that threaten the availability of critical systems. Similarly, CVE-2023-28231, a remote code execution vulnerability in Microsoft Windows DHCPv6 Server, and CVE-2024-38063, affecting Windows TCP/IP IPv6 implementations, have underscored the reality that even widely deployed, mature networking technologies contain exploitable flaws that could enable attackers to compromise critical infrastructure components. The National Security Agency has recognized these concerns, issuing specific guidance on IPv6 security considerations for critical infrastructure operators, acknowledging that the transition to next-generation networking protocols introduces new attack vectors that must be carefully managed.

The threat landscape extends far beyond individual software vulnerabilities to encompass systemic challenges inherent in how critical infrastructure networks are architected and operated. Research examining cybersecurity vulnerabilities in smart cities has identified multiple domains where security weaknesses persist, including energy systems, transportation networks, healthcare infrastructure, and municipal services [2]. These domains are increasingly interconnected, creating cascading vulnerability chains where a compromise in one system can propagate to others. The challenge is compounded by the heterogeneous nature of infrastructure networks, which often combine legacy systems designed without security considerations alongside modern IoT devices with varying levels of security maturity. This heterogeneity makes it exceptionally difficult to implement unified security policies or deploy consistent monitoring capabilities across the entire infrastructure ecosystem. Furthermore, the operational requirements of critical infrastructure impose constraints that do not exist in traditional IT environments; systems must maintain high availability, meet strict latency requirements, and operate continuously without the possibility of taking equipment offline for security updates or patches.

The threat actors targeting critical infrastructure range from opportunistic criminals to sophisticated nation-state adversaries, each bringing different capabilities and motivations to bear. Distributed denial of service attacks represents a persistent threat to IoT-enabled infrastructure, with research documenting comprehensive taxonomies of DDoS attack vectors specifically targeting connected devices [3]. These attacks can overwhelm network resources, disrupt service availability, and serve as diversions while more sophisticated intrusions occur. Beyond availability attacks, adversaries increasingly target the integrity and confidentiality of industrial control systems, seeking to manipulate operational parameters, steal proprietary information, or establish persistent access for future operations. The convergence of IoT technologies with critical infrastructure has created new attack surfaces that adversaries are actively exploiting. Studies have shown that IoT devices deployed in smart city environments frequently suffer from inadequate authentication mechanisms, unencrypted communications, and

insufficient input validation, creating opportunities for attackers to compromise these devices and use them as footholds for deeper network penetration [4], [5].

Table 1. Major Vulnerability Categories in Critical Infrastructure Networks

Vulnerability Category	Description	Example Manifestations	Potential Impact
Protocol Vulnerabilities	Flaws in fundamental networking protocols	CVE-2023-20080 (DHCPv6), CVE-2024-38063 (IPv6 TCP/IP)	Remote code execution, denial of service, network segmentation bypass
IoT Device Security	Weak authentication, unencrypted communications, inadequate updates	Default credentials, plaintext protocols, outdated firmware	Device compromise, botnet recruitment, network reconnaissance
Legacy System Integration	Systems designed without security considerations interfacing with modern networks	Unencrypted SCADA protocols, inadequate access controls	Operational disruption, data exfiltration, physical safety impacts
Supply Chain Risks	Compromised components or software in infrastructure systems	Malicious firmware, backdoored hardware	Long-term persistent access, widespread compromise
Configuration Weaknesses	Improper security settings and inadequate hardening	Open management interfaces, excessive permissions	Unauthorized access, privilege escalation

Sources: NSA (2023); Riggs et al. (2023); Haq et al. (2023)

3. MACHINE LEARNING APPROACHES TO INTRUSION DETECTION

Machine learning techniques for intrusion detection in critical infrastructure networks can be broadly categorized into supervised, unsupervised, and deep learning approaches, each with distinct strengths and limitations. Supervised learning methods rely on labeled datasets of normal and malicious traffic and can achieve high detection accuracy when trained on representative data. However, their effectiveness is limited in critical infrastructure environments where comprehensive labeled datasets are scarce, and attack patterns are often undocumented. Additionally, models trained on conventional IT network data may not generalize well to industrial control systems due to fundamental operational differences.

Recent reviews focused on critical infrastructure (CI) confirm these trade-offs in operational settings: ML-based IDS for ICS/SCADA/DCS must balance detection gains with scarce labeled data, zero-day variability, and resource limits that complicate real-time deployment and evaluation datasets. These surveys also compare modern CI datasets and highlight gaps when models trained on IT data are transferred to OT traffic. These findings reinforce the need for domain-aware model design and careful dataset curation in CI environments [24]

Unsupervised learning approaches address these limitations by modeling normal network behavior and detecting deviations as potential intrusions. This paradigm is particularly valuable for identifying zero-day attacks and novel threats in infrastructure networks. While anomaly detection techniques have shown promise in distinguishing malicious activity from normal operations, their practical deployment is challenged by high false positive rates. Legitimate operational variations caused by load fluctuations, maintenance activities, and seasonal effects can be difficult to differentiate from actual security incidents without highly nuanced models.

Deep learning methods, including convolutional and recurrent neural networks, offer advanced capabilities for detecting complex and temporally correlated attack patterns. These models can automatically extract features from high-dimensional network traffic data and are well suited for analyzing large-scale IoT and smart infrastructure environments. Despite their effectiveness, deep learning approaches present challenges related to high computational demands and limited interpretability, which raise concerns for deployment in resource-constrained and safety-critical infrastructure systems.

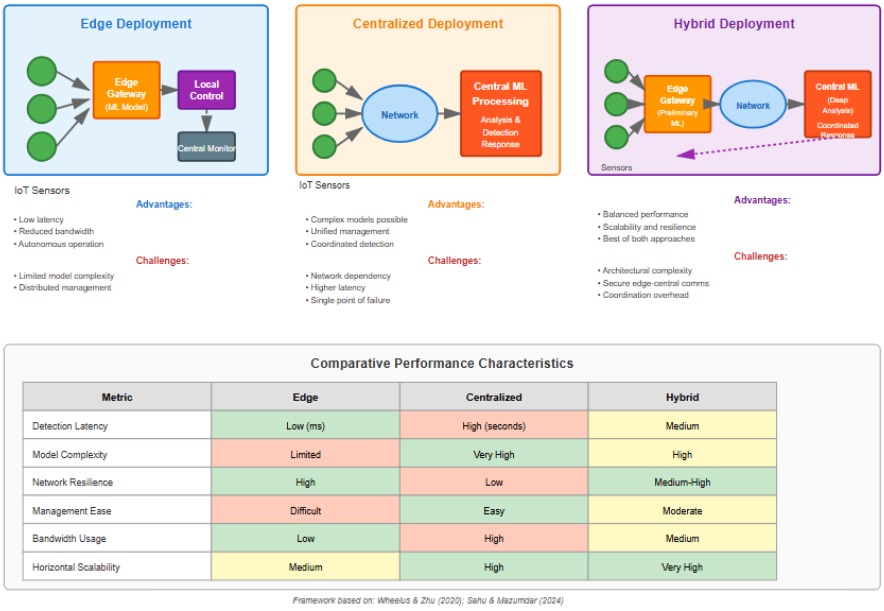

Metric	Edge	Centralized	Hybrid
Detection Latency	Low (ms)	High (seconds)	Medium
Model Complexity	Limited	Very High	High
Network Resilience	High	Low	Medium-High
Management Ease	Difficult	Easy	Moderate
Bandwidth Usage	Low	High	Medium
Horizontal Scalability	Medium	High	Very High

Figure 1. Machine Learning Model Deployment Architectures for Critical Infrastructure

The integration of machine learning into critical infrastructure intrusion detection systems must address fundamental questions about model deployment architecture and operational integration. Edge computing approaches, where machine learning models are deployed directly on network devices or local gateways, offer the advantage of reduced latency and decreased dependence on network connectivity to centralized processing resources. This architectural approach is particularly relevant for infrastructure systems with stringent real-time requirements or limited bandwidth availability. However, edge deployment constrains the complexity of models that can be employed, as resource-limited devices may lack the computational capacity to execute sophisticated deep learning architectures. Centralized approaches, conversely, enable the deployment of more complex models and facilitate coordination of detection across multiple infrastructure domains, but introduce latency that may be unacceptable for time-critical applications and create single points of failure. Hybrid architectures that combine edge-based preliminary analysis with centralized deep inspection represent a promising middle ground but add architectural complexity and introduce additional security considerations regarding the communication channels between edge and central components [6][4].

3.1. Technical Challenges in Infrastructure-Specific Intrusion Detection

The application of machine learning to critical infrastructure intrusion detection confronts technical challenges that distinguish this domain from conventional network security applications. One fundamental challenge concerns the imbalanced nature of security data in operational environments. Normal operational traffic vastly outnumbers malicious traffic in most

infrastructure networks, creating severe class imbalance problems that can cause machine learning models to achieve high overall accuracy while failing to detect actual intrusions. This imbalance is particularly acute in critical infrastructure contexts where attacks are relatively rare events, yet the consequences of failing to detect them can be catastrophic. Addressing class imbalance requires specialized techniques such as synthetic oversampling, cost-sensitive learning, or ensemble methods that explicitly account for the differential importance of correctly classifying minority class instances [7]. However, these techniques introduce their own complications, including the risk of overfitting to synthetic data or creating models that generate excessive false positives in attempts to maximize detection of rare attack patterns.

The concept drift phenomenon represents another significant technical challenge that has profound implications for the long-term effectiveness of machine learning-based intrusion detection systems. Network behavior in critical infrastructure environments is not static; it evolves over time due to factors including equipment upgrades, operational procedure changes, seasonal variations, and the introduction of new services. Machine learning models trained on historical data can become progressively less accurate as the underlying data distribution shifts away from the distribution on which they were trained. Research examining IoT security has identified concept drift as a critical concern for maintaining detection accuracy over extended operational periods [8]. Addressing concept drift requires implementing mechanisms for continuous model updating and retraining, but doing so in critical infrastructure environments raises challenging questions about how to validate updated models before deployment and how to ensure that model updates do not inadvertently introduce new vulnerabilities or degrade detection performance.A recent comprehensive survey of ML strategies for IDS further documents how model selection, metric choice, and evaluation protocols interact with drift and false-alarm management in production systems, underscoring the need for periodic recalibration and explainability to sustain trust [26]

Table 2. Technical Challenges in ML-Based Intrusion Detection for Critical Infrastructure

Challenge Domain	Specific Issues	Impact on Detection Systems	Mitigation Approaches
Data Imbalance	Attack traffic represents <0.1% of total traffic	High false negative rates, models biased toward majority class	Synthetic oversampling (SMOTE), cost-sensitive learning, ensemble methods
Concept Drift	Network behavior evolves over time	Degraded accuracy, increased false positives/negatives	Online learning, periodic retraining, ensemble adaptation
Feature Engineering	High dimensionality, protocol diversity, domain-specific semantics	Computational overhead, information loss, reduced interpretability	Domain expertise integration, automated feature selection, dimensionality reduction
Adversarial Robustness	Poisoning attacks, evasion techniques, model exploitation	Compromised detection capability, false confidence	Adversarial training, robust optimization, input sanitization
Real-time Processing	Latency constraints, throughput requirements	Delayed detection, incomplete analysis	Edge computing, lightweight models, hierarchical processing

Sources: Kumari & Jain (2023); Noor & Hassan (2019); Comprehensive Study (2025)

4. FEATURE ENGINEERING, ADVERSARIAL RISKS, AND OPERATIONAL CHALLENGES

Feature engineering in critical infrastructure networks is challenging due to diverse industrial protocols and complex operational behaviors that generate high-dimensional feature spaces. Unlike traditional IT environments, these systems embed semantic information tied to physical processes, requiring domain expertise to extract security-relevant features. At the same time, models must balance feature richness with computational efficiency, as overly complex representations increase latency and hinder real-time detection.Hybrid ML/DL pipelines that combine ensemble learners with deep architectures have recently shown strong performance while explicitly managing class imbalance and feature reduction to keep inference latency acceptable [27]. Such designs illustrate a practical path to reconcile feature richness with real-time requirements in modern NIDS [27].

Adversarial machine learning further complicates intrusion detection. Attackers can poison training data, craft evasive traffic, or exploit model weaknesses, turning ML-based detectors into targets themselves. Robust training, adversarial defenses, and continuous validation are essential but often add computational overhead that conflicts with strict real-time constraints in critical infrastructure.

Operational deployment introduces additional barriers. Infrastructure systems require near-continuous availability, leaving little flexibility for model retraining or system downtime. Insufficiently validated model updates risk false negatives that enable attacks or false positives that burden security teams. Effective integration therefore demands structured alert prioritization, workflow alignment, and human-centered oversight, supported by explainable models that allow analysts to interpret detection outputs.

Regulatory and workforce limitations also hinder adoption. Frameworks such as NERC CIP and the America's Water Infrastructure Act provide limited guidance for validating ML-based security systems, creating uncertainty around compliance. Simultaneously, a shortage of professionals with combined expertise in cybersecurity, machine learning, and operational technology constrains many organizations—particularly smaller operators—from deploying and sustaining advanced intrusion detection capabilities.

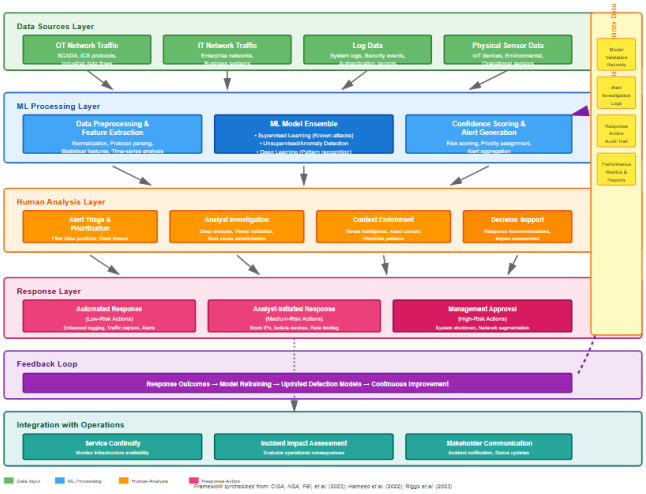

Figure 2. Integration Framework for ML-Based Intrusion Detection in Infrastructure Operations

4.1. Smart Grid and Power Infrastructure Security

The electric power grid remains a core critical infrastructure where machine-learning-based intrusion detection can significantly strengthen security, yet it also highlights the operational challenges of protecting industrial control environments. Smart grid modernization has introduced extensive digital communication and control capabilities across generation, transmission, and distribution systems. While these advances improve efficiency and reliability, they also expand the cyber attack surface. Research identifies multiple attack vectors—including threats to SCADA systems, advanced metering infrastructure, distributed energy resources, and grid communication networks—each capable of causing disruptions ranging from localized outages to large-scale cascading failures [9]. ML-driven anomaly detection offers promising early-warning capabilities, but deployment must account for the grid's strict real-time and safety-critical requirements.

Modern smart grids operate through layered communication networks spanning high-voltage transmission systems, distribution networks, and customer-level infrastructure. Each tier exhibits unique protocols, data flows, and timing behavior, making it difficult for a single intrusion detection model to generalize across the entire grid. Studies emphasize that effective protection requires domain-specific ML models tuned to each architectural layer, though this specialization increases operational complexity and resource demands [9].

False positives pose particularly high risks in the power sector. Unlike IT networks, where isolating suspicious activity typically has limited physical consequences, automated responses in grid environments can trigger protective relays, disconnect generation assets, or unintentionally interrupt service. This creates an inherent tension between rapid automated mitigation and the need for human oversight in safety-critical scenarios. Accordingly, response architectures must differentiate between low-risk automated actions (e.g., increased logging, deeper traffic inspection) and high-risk control actions that require operator approval. ML systems must therefore deliver interpretable outputs, providing clear confidence indicators and contextual information that enable operators to make informed, safe decisions during anomalous events.

Table 3. Machine Learning Applications in Smart Grid Security

Grid Domain	Security Challenges	ML Application	Expected Benefits	Implementation Challenges
SCADA Systems	Unauthorized control commands, data manipulation	Supervised learning for command validation	Detection of anomalous control sequences	Limited attack training data, real-time requirements
Advanced Metering Infrastructure	Meter tampering, data privacy, communication attacks	Anomaly detection on meter data patterns	Identification of compromised meters	Massive data volumes, legitimate variability
Distributed Energy Resources	Inverter manipulation, coordinated attacks	Time-series analysis of resource behavior	Detection of coordinated anomalies	Weather-driven variations, heterogeneous devices
Grid Communication Networks	Protocol attacks, eavesdropping, replay attacks	Deep learning on network traffic	Protocol anomaly detection	Specialized OT protocols, latency constraints

Substation Automation	IEC 61850 protocol attacks, device compromise	Ensemble methods combining multiple detection approaches	Comprehensive threat detection	Integration with legacy systems

5. MUNICIPAL SYSTEMS AND SMART CITY SECURITY

Municipal smart city implementations present a particularly complex challenge for machine learning-based intrusion detection due to the diversity of systems involved and the public safety implications of security failures. Smart cities integrate numerous infrastructure domains including traffic management, public safety communications, environmental monitoring, waste management, and municipal services into interconnected digital ecosystems. Research examining smart city security has documented the wide-ranging cyber threats facing these systems, from attacks targeting individual IoT sensors to sophisticated campaigns aiming to disrupt critical municipal services [10]. The interconnected nature of smart city systems means that a compromise in one domain can potentially cascade to affect other municipal services, creating compound effects that could significantly impact public safety and quality of life. Machine learning-based intrusion detection deployed across smart city infrastructure must therefore account for the possibility of cross-domain attacks and lateral movement between different municipal systems.

The privacy implications of security monitoring in smart city environments warrant careful consideration, as the data collected for intrusion detection purposes may contain information about citizen behavior and activities. Traffic monitoring systems designed to detect anomalies that could indicate cyber-attacks will necessarily collect information about vehicle movements throughout the city. Environmental sensors deployed for pollution monitoring could reveal patterns of human activity. Research has emphasized that smart city implementations must carefully balance security monitoring requirements against citizen privacy rights, implementing appropriate data minimization, anonymization, and retention policies [11]. Machine learning systems must be designed to extract security-relevant features from municipal network traffic while avoiding unnecessary collection or retention of personally identifiable information. This requires thoughtful feature engineering that focuses on network-level characteristics rather than content, and implementation of data governance frameworks that specify clear policies for how security monitoring data will be handled.

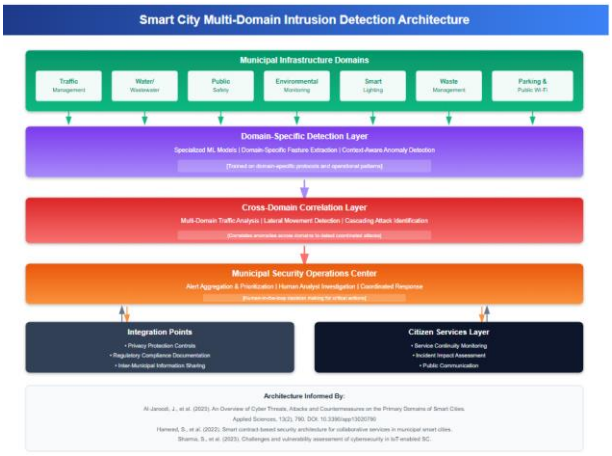

Figure 3. Smart City Multi-Domain Intrusion Detection Architecture

The heterogeneity of devices and systems deployed in smart city environments creates significant challenges for developing comprehensive intrusion detection capabilities. A single municipal network may encompass thousands of different device types from dozens of vendors, operating various protocols and exhibiting diverse behavioral characteristics. Research examining IoT vulnerabilities and countermeasures has highlighted how this heterogeneity complicates efforts to establish baseline normal behavior and detect anomalies [12]. Machine learning approaches must either develop device-specific models tailored to each class of equipment, requiring extensive training data and ongoing maintenance for each device type, or develop more general models that can accommodate diverse device behaviors while still maintaining adequate detection sensitivity. Neither approach is entirely satisfactory; device-specific models impose substantial operational overhead, while general models may sacrifice detection accuracy in pursuit of broader applicability. Hybrid approaches that group devices into behavioral clusters and develops specialized models for each cluster represent a promising middle ground but require sophisticated clustering algorithms and ongoing model management.

5.1. Data Management and Privacy Considerations

The effectiveness of machine-learning-based intrusion detection in critical infrastructure depends heavily on high-quality training data, yet such data is difficult to obtain. Unlike conventional IT environments with abundant public intrusion datasets, infrastructure-specific data is scarce due to the sensitivity of operational technology networks and concerns that sharing traffic captures or logs could expose vulnerabilities. Studies highlight that limited access to representative datasets remains a major barrier to developing tailored ML-based security solutions for IoT and OT environments [4]. As a result, researchers and operators face persistent data scarcity that restricts robust model development and evaluation.

Synthetic data generation—using simulations or generative models—offers a potential workaround by producing artificial traffic resembling real operational behavior. While useful for prototyping, synthetic datasets often struggle to reproduce the full complexity, edge-case interactions, and unexpected anomalies found in real systems. Consequently, questions remain about whether models trained solely on synthetic traffic can generalize reliably in production environments, where validation itself may require access to sensitive operational data.

Data retention and governance add further complexity. Effective ML models benefit from long-term historical datasets capturing seasonal variations and diverse operating conditions. However, retaining large volumes of network traffic and security logs increases the risk of exposing sensitive infrastructure information in the event of a breach. Prior research emphasizes that strong data governance—defining what data is collected, how it is protected, retention periods, and permitted sharing conditions—is essential to balancing security, regulatory, and privacy requirements [13]. Developing these policies requires coordination with legal, privacy, and regulatory stakeholders to ensure compliance while maintaining the data pipeline necessary for reliable machine learning–based intrusion detection.

Table 4. Data Challenges and Solutions for Infrastructure ML Security

Data Challenge	**Security Impact**	**Potential Solutions**	**Trade-offs**
Training Data Scarcity	Limited model effectiveness, poor generalization	Synthetic data generation, simulation environments, data sharing consortia	Synthetic data may not capture real-world complexity, sharing faces confidentiality concerns
Dataset Imbalance	Biased models, poor attack detection	Oversampling techniques, cost-sensitive learning,	Risk of overfitting, computational overhead

		transfer learning	
Privacy Constraints	Limited data sharing, reduced collaborative learning	Federated learning, differential privacy, secure multi-party computation	Performance penalties, implementation complexity
Data Quality Issues	Inaccurate labels, incomplete captures, noise	Automated quality checking, human validation, robust learning algorithms	Resource intensive validation, cannot eliminate all errors
Temporal Coverage	Models trained on limited time periods	Long-term data collection, seasonal dataset augmentation	Storage costs, evolving attack landscape may date historical data

6. OPPORTUNITIES AND EMERGING DIRECTIONS

Despite the challenges of applying machine learning to critical infrastructure intrusion detection, several promising opportunities can meaningfully advance security capabilities. One key direction is the development of ensemble-based intrusion detection systems that combine supervised, unsupervised, and deep learning methods. Such ensembles leverage the complementary strengths of multiple approaches, reducing false positives while improving detection sensitivity compared to single-model solutions [3]. The primary challenge is designing robust fusion mechanisms that appropriately weight model outputs and resolve conflicts when classifications diverge.Recent empirical results show that hybrid ensembles that pair tree-based feature selection (e.g., XGBoost) with CNN/LSTM temporal learners can reduce false alarms while improving detection across CICIDS2017, UNSW-NB15, and NSL-KDD datasets, strengthening the practical case for fusion architectures [28]

Explainable artificial intelligence (XAI) presents another important opportunity. Recent advances—such as attention mechanisms that highlight influential features and counterfactual explanations that show how inputs must change to alter outcomes—offer potential for increasing transparency and operator trust [15], [16]. For intrusion detection, XAI can help analysts understand why traffic was flagged as malicious and support more effective investigation. The remaining challenge lies in tailoring explanations to operational security workflows so they are actionable and not overly complex.

Transfer learning and domain adaptation also hold promise for addressing training data scarcity. Rather than requiring every operator to build models from scratch, transfer learning enables models trained in one infrastructure domain to be adapted to another with limited fine-tuning [8], [17]. The key research question is determining which model components generalize across domains and which must be customized for infrastructure-specific protocols and operational patterns.

Physics-informed machine learning represents another promising direction. By incorporating domain knowledge about physical system behavior into learning processes—such as hydraulic relationships in water systems or power flow dynamics in electrical grids—models can better distinguish between benign operational anomalies and true cyber threats. This approach enables detection that aligns more closely with real-world system behavior, reducing false positives and improving operator confidence [18]. Achieving this requires sustained collaboration between ML researchers, cybersecurity experts, and domain specialists.

Testbed environments offer additional opportunities for advancing ML-based intrusion detection. Platforms such as OpenCyberCity provide realistic, controlled environments for evaluating algorithms, validating robustness, and enabling adversarial testing prior to deployment [19]. These environments also support synthetic dataset generation under realistic conditions, helping mitigate data scarcity. However, testbeds must evolve to reflect the complexity and unpredictability of real infrastructure; strong performance in a laboratory environment does not guarantee operational effectiveness.

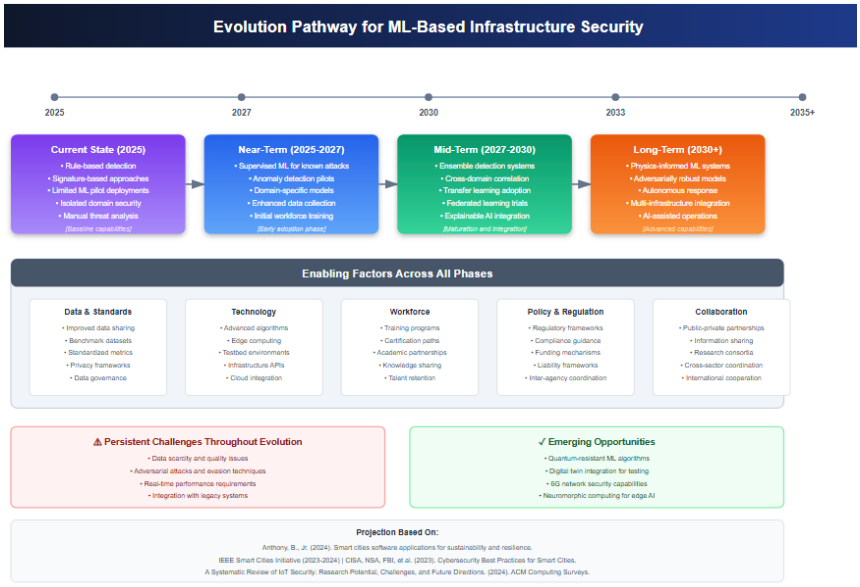

Figure 4. Evolution Pathway for ML-Based Infrastructure Security

7. CROSS-SECTOR COLLABORATION AND INFORMATION SHARING

The development of effective machine learning-based intrusion detection for critical infrastructure cannot succeed as isolated efforts by individual organizations, but rather requires coordinated collaboration across sectors, sharing of threat intelligence, and collective advancement of security capabilities. The challenges facing different infrastructure sectors share common characteristics, and lessons learned in one domain can inform security approaches in others. Research has emphasized the importance of collaborative approaches to smart city cybersecurity, noting that municipalities can benefit substantially from sharing security intelligence and coordinating defensive strategies [20]. However, establishing effective information sharing mechanisms faces obstacles including competitive concerns, liability questions, and the technical challenges of sharing actionable security information while protecting sensitive operational details.

The establishment of Information Sharing and Analysis Centers specific to various infrastructure sectors has created formal mechanisms for coordinating security efforts, but these organizations face ongoing challenges in facilitating effective information exchange about emerging threats and effective countermeasures. Machine learning-based intrusion detection introduces new dimensions to information sharing, as organizations could potentially benefit from sharing not just threat indicators but also detection models, feature engineering approaches, and training datasets. Research examining cybersecurity in industrial IoT has explored how blockchain technologies could potentially facilitate trusted information sharing while maintaining auditability and attribution [21]. The development of privacy-preserving mechanisms for sharing

machine learning models and training data could enable infrastructure operators to collectively develop improved security capabilities without exposing proprietary operational information [22].

Table 5. Stakeholder Roles in ML-Based Infrastructure Security Development

Stakeholder Category	Primary Responsibilities	Contribution to ML Security	Critical Success Factors
Infrastructure Operators	Deploy and operate security systems, respond to incidents	Provide operational context, validate solutions, share anonymized data	Executive commitment, adequate resources, workforce development
Technology Vendors	Develop security products, provide technical support	Integrate ML capabilities, leverage cross-customer insights	Balance proprietary concerns with collaboration, long-term support commitment
Government Agencies	Set standards, provide guidance, facilitate information sharing	Develop regulatory frameworks, fund research, coordinate sector collaboration	Clear policy direction, adequate funding, cross-agency coordination
Academic Researchers	Conduct fundamental research, develop novel approaches	Advance ML techniques, evaluate approaches, train workforce	Access to representative data, engagement with practitioners, research funding
Industry Associations	Facilitate information sharing, develop best practices	Coordinate collaborative security initiatives, aggregate sector knowledge	Member engagement, trusted neutral platform, technical expertise
Security Service Providers	Deliver managed security services, incident response	Scale expertise across multiple operators, provide specialized capabilities	Access to threat intelligence, trained analysts, technology integration

The role of equipment vendors and technology providers in supporting machine learning-based security deserves particular attention, as these organizations have visibility across multiple customer deployments and could potentially identify security patterns that individual operators cannot observe. Vendors of industrial control systems, IoT devices, and infrastructure management platforms accumulate experience from numerous installations and could leverage this collective experience to develop security capabilities that benefit all customers. Research has explored how vendor-provided security services could incorporate machine learning capabilities that are continuously improved based on anonymized telemetry from customer deployments (Research on IPv6 Address State Detection and Management Technology, 2024). However, this model raises questions about data ownership, privacy protection, and the potential for vendor lock-in, requiring careful contractual and technical safeguards to protect customer interests while enabling beneficial security collaboration.

Academic research institutions have important roles to play in advancing the state of practice for infrastructure security, conducting fundamental research on machine learning techniques, evaluating proposed approaches, and training the next generation of security professionals. The translation of academic research into operational practice faces well-documented challenges, as research environments typically do not capture the full complexity and operational constraints of production infrastructure. Strengthening partnerships between academic researchers and

infrastructure operators can help ensure that research addresses practically relevant problems and that promising research results are effectively transitioned to operational deployment. Research examining security vulnerability assessments has noted the importance of diverse research communities bringing different perspectives to cybersecurity challenges [23]. Supporting collaborative research that brings together computer scientists, infrastructure engineers, and security practitioners can produce solutions that are both technically sophisticated and operationally viable.

8. STANDARDIZATION AND EVALUATION FRAMEWORKS

Advancing machine-learning-based intrusion detection for critical infrastructure requires standardized evaluation frameworks that allow objective comparison of techniques and give operators confidence in security performance. Although numerous ML approaches report high accuracy, these results often rely on different datasets, metrics, and threat models, making cross-study comparison difficult. Prior research emphasizes that the absence of uniform evaluation methodologies significantly hinders assessment of competing IoT and OT security solutions (A review of the security vulnerabilities and countermeasures in the Internet of Things solutions, 2023). Establishing shared benchmarks, metrics, and testing procedures would accelerate adoption by enabling operators to make evidence-based technology decisions.

A key requirement is the development of benchmark datasets that accurately represent operational technology environments. Existing public datasets focus on conventional IT networks and do not capture the protocols, traffic patterns, or attack vectors seen in industrial systems. Although emerging research has begun creating infrastructure-specific datasets, comprehensive multi-domain benchmarks remain limited. Synthetic datasets—generated through simulations or sanitized real traffic—offer partial solutions but must be rigorously validated to ensure they reflect the complexity and behavioral nuances essential for evaluating intrusion detection performance.

Standardizing performance metrics is equally important. Infrastructure environments impose constraints that conventional measures such as precision, recall, and F-score do not fully capture. The asymmetric costs of errors—where a single false positive may trigger an unnecessary shutdown or misoperation—demand metrics that incorporate safety impact, operational disruption, and latency requirements. High detection accuracy is insufficient if achieved at the expense of real-time responsiveness or explainability. Recent studies highlight the need for holistic evaluation frameworks that integrate detection performance with operational impact, resource demands, and implementation complexity [21].

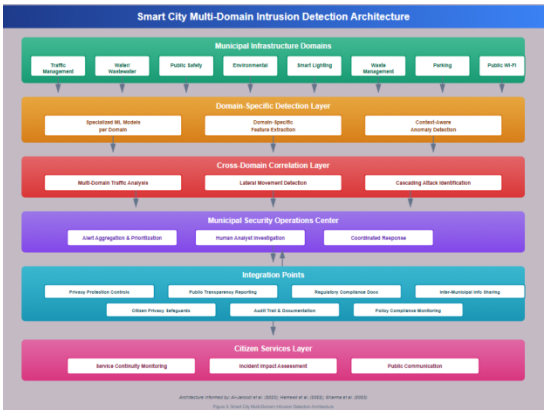

Figure 5. Comprehensive Evaluation Framework for Infrastructure ML Security Systems

The certification and validation of machine learning-based security systems for use in safety-critical infrastructure applications presents significant challenges that the security community has only begun to address. Traditional software assurance approaches based on code review, formal verification, and exhaustive testing do not translate straightforwardly to machine learning systems whose behavior is determined by training data rather than explicit programming. How can an infrastructure operator gain confidence that a machine learning intrusion detection system will perform reliably across the full range of operational conditions it might encounter? What testing and validation procedures provide adequate assurance that the system will not produce dangerous false positives or miss critical attacks? Research examining AI assurance has explored various approaches including adversarial testing, model interpretability analysis, and formal verification of neural network properties, but significant work remains to translate these research concepts into practical validation procedures appropriate for infrastructure security applications. The development of industry consensus around appropriate validation approaches could help accelerate the adoption of machine learning security technologies by providing operators with confidence that deployed systems meet appropriate assurance standards.Recent IDS surveys further emphasize that progress depends on reproducible pipelines using modern datasets such as CIC-IDS2017, CICDDoS2019, and UNSW-NB15, along with metrics that go beyond accuracy (e.g., precision/recall balance, false-positive rate, and detection latency), to support meaningful cross-study comparison [25]

9. FUTURE RESEARCH DIRECTIONS AND RECOMMENDATIONS

Advancing machine-learning-based intrusion detection for critical infrastructure requires coordinated progress across research, technology development, and operational practice. Several priority directions emerge from current challenges. First, new machine learning methods must be designed specifically for the characteristics of operational technology networks. Many existing approaches rely on general-purpose algorithms not built for constraints such as extreme data imbalance, concept drift, adversarial manipulation, and explainability requirements. Research that integrates these needs into unified algorithmic frameworks could produce models far better aligned with real-world infrastructure environments.

Second, the intersection of machine learning and formal methods represents an important but underdeveloped research frontier. Formal verification offers mathematically grounded assurances about system behavior—critical for safety-sensitive environments—but current progress remains limited to relatively small neural network architectures. Extending verification techniques to larger, more complex models and validating properties such as adversarial robustness or bounded false-positive rates would help address assurance concerns that currently hinder adoption.

Third, human-AI collaboration frameworks tailored to infrastructure security operations warrant focused attention. Optimal security practice will likely blend automated detection with human expertise rather than rely exclusively on either. Research is needed to design interfaces, workflows, and decision-support tools that empower analysts to interpret model outputs, understand uncertainty, and intervene appropriately under real-world operational constraints. Insights from human–computer interaction, cognitive psychology, and security operations research will be crucial for building effective collaboration models.

From a policy and governance perspective, government agencies can accelerate progress by funding initiatives dedicated to developing and validating ML-based security tools for critical infrastructure. Such programs could support benchmark datasets, reference implementations, and standardized validation frameworks. Regulatory bodies should also update compliance guidelines to address machine-learning-based detection systems explicitly, reducing uncertainty for operators evaluating deployment.

Infrastructure operators, in turn, should prepare for increased reliance on ML-driven security by investing in data collection infrastructure, workforce upskilling, and operational processes that enable effective model deployment. Proactive engagement with researchers and technology vendors—such as participation in pilot deployments and field trials—can help organizations adopt emerging solutions safely and gain strategic advantages.

Finally, the academic research community must balance scientific innovation with practical relevance. Sustained collaboration with infrastructure operators is essential to ensure models address genuine operational constraints and produce solutions that remain usable outside laboratory conditions. Research should both advance generalizable knowledge and deliver approaches that can be realistically integrated into production environments. The most impactful work will bridge theory and practice, addressing immediate operational needs while contributing to long-term advancements in securing critical infrastructure.

10. CONCLUSION

Machine learning offers significant potential to enhance intrusion detection across critical infrastructure by identifying sophisticated attacks, detecting previously unseen threats, and adapting to evolving adversary behaviors. However, translating this potential into operational reality remains challenging due to limitations in algorithm design, data availability, system integration, regulatory uncertainty, and workforce preparedness. Although recent research has advanced understanding and proposed promising solutions, considerable work is still required to achieve reliable, scalable deployment in real-world infrastructure environments.

Realizing the benefits of machine learning will require coordinated action across stakeholder groups—including infrastructure operators, technology vendors, government agencies, academic researchers, and industry associations. No single sector can address the multifaceted technical and organizational barriers alone. Progress depends on sustained information sharing, collaborative research, and the development of standardized evaluation frameworks, benchmark datasets, and validation procedures tailored to safety-critical systems. Workforce development will also be essential to ensure organizations can effectively deploy, manage, and oversee advanced ML-based security technologies.

As critical infrastructure becomes more interconnected and dependent on digital controls, traditional security approaches such as network segmentation and signature-based detection alone are increasingly insufficient. Machine learning provides a pathway to scalable, adaptive, and real-time security capabilities capable of handling the complexity and volume of modern IoT-enabled environments. Achieving this vision requires long-term commitment to overcoming the technical and operational challenges outlined throughout this work. The research community has laid a strong foundation, and continued innovation will drive progress toward mature and trustworthy machine-learning-based security solutions over the coming decade.

Moving forward, stakeholders must balance optimism with practical considerations. Operators should avoid both premature adoption of immature technologies and excessive caution that delays meaningful improvements. Vendors must prioritize solutions aligned with operational constraints rather than focusing solely on laboratory performance. Researchers must stay grounded in real deployment needs, while policymakers must develop regulatory frameworks that encourage innovation while ensuring safety and accountability. Through coordinated and sustained effort, machine-learning-enhanced security for U.S. critical infrastructure can evolve from aspiration to reality, providing the protection necessary for an increasingly digital society.

10. ACKNOWLEDGEMENTS

The authors would like to express their sincere appreciation to colleagues, collaborators, and mentors whose technical insights, reviews, and constructive discussions contributed to the development of this research. The authors also acknowledge the support of professional and academic environments that enabled access to domain knowledge, analytical tools, and research perspectives essential to the study of machine learning–driven intrusion detection in critical infrastructure systems.

REFERENCES

[1] I. Rafiq, A. Mahmood, S. Razzaq, S. H. M. Jafri, and I. Aziz, "IoT applications and challenges in smart cities and services," *The Journal of Engineering*, vol. 2023, no. 4, pp. 1–20, Mar. 2023, doi: 10.1049/tje2.12262.

[2] V. Demertzi, S. Demertzis, and K. Demertzis, "An overview of cyber threats, attacks and countermeasures on the primary domains of smart cities," *Applied Sciences*, vol. 13, no. 2, p. 790, Jan. 2023, doi: 10.3390/app13020790.

[3] P. Kumari and A. K. Jain, "A comprehensive study of DDoS attacks over IoT network and their countermeasures," *Computers & Security*, vol. 127, p. 103096, Apr. 2023, doi: 10.1016/j.cose.2023.103096.

[4] S. K. Sahu and K. Mazumdar, "Exploring security threats and solutions techniques for Internet of Things (IoT): from vulnerabilities to vigilance," *Frontiers in Artificial Intelligence*, vol. 7, May 2024, Art. no. 1397480, doi: 10.3389/frai.2024.1397480.

[5] M. Farhan, H. Waheed ud din, S. Ullah, M. S. Hussain, M. A. Khan, T. Mazhar, U. F. Khattak, and I. H. Jaghdam, "Network-based intrusion detection using deep learning technique," *Scientific Reports*, vol. 15, no. 1, p. 8770, Jul. 2025, doi: 10.1038/s41598-025-08770-0.

[6] C. Wheelus and X. Zhu, "IoT network security: threats, risks, and a data-driven defense framework," *IoT*, vol. 1, no. 2, pp. 259–285, Oct. 2020, doi: 10.3390/iot1020016.

[7] B. Susilo, A. Muis, and R. F. Sari, "Intelligent intrusion detection system against various attacks based on a hybrid deep learning algorithm," *Sensors*, vol. 25, no. 2, p. 580, Jan. 2025, doi: 10.3390/s25020580.

[8] B. K. Mohanta, D. Jena, U. Satapathy, and S. Patnaik, "Survey on IoT security: Challenges and solution using machine learning, artificial intelligence and blockchain technology," *Internet of Things*, vol. 11, p. 100227, Sep. 2020, doi: 10.1016/j.iot.2020.100227.

[9] T. Mazhar, H. M. Irfan, S. Khan, I. Haq, I. Ullah, M. Iqbal, and H. Hamam, "Analysis of cyber security attacks and its solutions for the smart grid using machine learning and blockchain methods," *Future Internet*, vol. 15, no. 2, p. 83, Feb. 2023, doi: 10.3390/fi15020083.

[10] J. Telo, "Smart city security threats and countermeasures in the context of emerging technologies," *International Journal of Intelligent Automation and Computing*, vol. 6, no. 1, pp. 31–45, 2023.

[11] E. Ismagilova, L. Hughes, N. P. Rana, and Y. K. Dwivedi, "Security, privacy and risks within smart cities: literature review and development of a smart city interaction framework," *Information Systems Frontiers*, vol. 24, pp. 393–414, 2022, doi: 10.1007/s10796-020-10044-1

[12] M. Domb and Y. Shnaps, "Cybersecurity threats and mitigations related to smart cities operation," in *Smart Cities – Foundations and Perspectives*. London, U.K.: IntechOpen, 2024, doi: 10.5772/intechopen.114926

[13] M. B. M. Noor and W. H. Hassan, "Current research on internet of things (IoT) security: a survey," *Computer Networks*, vol. 148, pp. 283–294, Jan. 2019, doi: 10.1016/j.comnet.2018.11.025.

[14] P. Kumari and A. K. Jain, "A comprehensive study of DDoS attacks over IoT network and their countermeasures," *Computers & Security*, vol. 127, p. 103096, Apr. 2023, doi: 10.1016/j.cose.2023.103096.

[15] V. Z. Mohale and I. C. Obagbuwa, "A systematic review on the integration of explainable artificial intelligence in intrusion detection systems to enhancing transparency and interpretability in cybersecurity," *Frontiers in Artificial Intelligence*, vol. 8, 2025, Art. no. 1526221, doi: 10.3389/frai.2025.1526221.

[16] P. Barnard, N. Marchang, N. Khanna, and J. J. van Vuuren, "Decision tree-based machine learning models for cyber security intrusion detection," *International Journal of Information Security*, vol. 23, pp. 145–168, 2024, doi: 10.1007/s10207-023-00732-1.

[17] J. Zhang, C. Luo, C. Marcus, and G. Min, "Federated learning for distributed IIoT intrusion detection using transfer approaches," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 12, pp. 8450–8459, Dec. 2021, doi: 10.1109/TII.2021.3070617.

[18] A. Ullah, S. M. Anwar, J. Li, L. Nadeem, T. Mahmood, A. Rehman, and T. Saba, "Smart cities: the role of Internet of Things and machine learning in realizing a data-centric smart environment," *Complex & Intelligent Systems*, vol. 10, pp. 1607–1637, 2024, doi: 10.1007/s40747-023-01175-4

[19] A. (OpenCyberCity Testbed Team), "Development of the OpenCyberCity Testbed: Smart City Research Innovation and Opportunities," in *Proc. ACM/IEEE 14th Int. Conf. on Cyber-Physical Systems (ICCPS)*, 2023, doi: 10.1145/3576841.3589612.

[20] U.S. CISA, NSA, FBI, et al., "Cybersecurity Best Practices for Smart Cities," Joint Cybersecurity Advisory, 2023. *(Government advisory—cite with access date if required.)*

[21] A. Muneer, S. M. Fati, and S. Fuddah, "A critical review of artificial intelligence-based approaches in intrusion detection: A comprehensive analysis," *Journal of Engineering*, 2024, Art. ID 3909173, doi: 10.1155/2024/3909173.

[22] Y. Chen, X. Li, J. Wang, and L. Zhang, "Enhancing network security: leveraging machine learning for integrated protection and intrusion detection," *Intelligent Automation & Soft Computing*, vol. 40, no. 1, pp. 1–27, 2025, doi: 10.32604/iasc.2024.058624.

[23] A. Ghaffarian and M. Shahriari, "Research communities in cyber security vulnerability assessments: A comprehensive literature review," *Computers & Security*, vol. 126, p. 103187, 2023, doi: 10.1016/j.cose.2022.103187

[24] A. Pinto, L.-C. Herrera, Y. Donoso, and J. A. Gutierrez, "Survey on intrusion detection systems based on machine learning techniques for the protection of critical infrastructure," *Sensors*, vol. 23, no. 5, p. 2415, Feb. 2023, doi: 10.3390/s23052415.

[25] A. Hozouri, A. Mirzaei, and M. Effatparvar, "A comprehensive survey on intrusion detection systems with advances in machine learning, deep learning and emerging cybersecurity challenges," *Discover Artificial Intelligence*, vol. 5, Art. no. 314, Nov. 2025, doi: 10.1007/s44163-025-00578-1.

[26] A. Hussein Ali, M. Charfeddine, B. Ammar, *et al.*, "Unveiling machine learning strategies and considerations in intrusion detection systems: a comprehensive survey," *Frontiers in Computer Science*, vol. 6, 2024, Art. no. 1387354, doi: 10.3389/fcomp.2024.1387354.

[27] A. Bouzaachane, E. M. El Guarmah, A. M. Alnajim, and S. Khan, "Addressing modern cybersecurity challenges: a hybrid machine learning and deep learning approach for network intrusion detection," *Computers, Materials & Continua*, vol. 84, no. 2, pp. 2391–2410, 2025, doi: **10.32604/cmc.2025.065031**.

[28] M. Sajid, K. R. Malik, A. Almogren, *et al.*, "Enhancing intrusion detection: a hybrid machine and deep learning approach," *Journal of Cloud Computing*, vol. 13, Art. no. 123, Jul. 2024, doi: 10.1186/s13677-024-00685-x.

AUTHORS

Joy Selasi Agbesi is a senior Optical Network Engineer, Network Consultant, and researcher with over a decade of experience spanning telecom operators, global vendors, and hyperscale data-center environments. He has held technical and leadership roles at Huawei Technologies, MTN Ghana, and Meta, where he has led large-scale Optical Transport Network (OTN/DWDM) and IP/MPLS backbone projects, including nationwide modernizations and high-capacity (100G/400G) expansions supporting millions of subscribers across West Africa. His work bridges deep hands-on engineering with strategic network planning, capacity provisioning, and operational excellence.

Beyond optical networking, Joy brings strong interdisciplinary expertise in cybersecurity, artificial intelligence, and machine learning, applying intelligent and data-driven approaches to network security, anomaly detection, and resilient infrastructure design. He holds a BSc in Telecommunications Engineering from KNUST (Ghana) and an MSc in Information & Telecommunication Systems from Ohio University (USA), is a Senior Member of IEEE, and actively serves as a peer reviewer and technical program committee member for leading IEEE and ACM venues. As an author of multiple peer-reviewed

publications, his work integrates optical networking, security, and AI-driven systems to advance scalable, secure, and future-ready digital infrastructures.

Abigail Nanayaa Otchillis a Network Engineer at Meta with deep expertise in hyperscale data-center networking, AI infrastructure, and cybersecurity. She supports end-to-end rack orchestration across Meta's global data centers, leading network provisioning, optical terminal deployments, capacity expansions, and router and switch upgrades that power Facebook, Instagram, data platforms, and large-scale AI/ML workloads. Abigail also leads continuous data-center infrastructure improvements, optimizing network performance, reliability, and uptime to support mission-critical services at scale. Complementing her strong technical foundation, Abigail holds CIPM and CISA certifications and is an IAPP Westin Scholar, bringing a unique blend of cybersecurity, privacy engineering, and AI-aware network operations to the design and operation of resilient, high-performance data-center environments.

Raymond Horlalie Tay, is a Senior Optical Network Engineer at Comcast, with over 14 years of experience architecting mission-critical infrastructure. Previously, he served in the same capacity at Amazon Web Services (AWS), where he validated and implemented 40+ Tb/s optical transport networks specifically designed to support hyperscale AI-driven cloud workloads. His expertise extends to the intersection of cybersecurity and critical infrastructure; at Convergent Energy + Power, he led the deployment of NIST 800-compliant secure SD-WAN architectures to protect national-scale renewable energy assets spanning the US, Canada and Puerto Rico. His research interests focus on the intersection of photonic network scalability, machine learning, and artificial intelligence, particularly in optimizing flex-grid optical layers to handle the bursty, high-bandwidth demands of distributed machine learning training, while ensuring the security and resilience of the underlying physical transport. He holds a Master of Science in Telecommunication Networks from Northeastern University, Boston.

Noah K. Bamfo is a cybersecurity and network engineering professional with over nine years of experience designing secure, intelligent systems that leverage artificial intelligence and machine learning to enhance network defenses, particularly at the DNS layer. He holds advanced degrees in computer science and telecommunications and has led high-impact security deployments for government, financial, and enterprise institutions in Africa and the United States.