Business Data Analytics for Cyber Risk Detection in Distributed IoT-Enabled Operations

Andika Priyono[1], Rini Kurnia Sari[2], Bambang Susilo[3,*]

[1]School of Information Systems, Universitas Mercu Buana, Jakarta 11650, Indonesia

[2]Department of Industrial Engineering, Universitas Tarumanagara, Jakarta 11440, Indonesia

[3]Department of Business Analytics, Universitas Atma Jaya Yogyakarta, Yogyakarta 55281, Indonesia

*Email: bambang.susilo@uajy.ac.id (Corresponding Author)

Abstract

The proliferation of Internet of Things (IoT) devices across manufacturing, healthcare, energy, and mobility operations has expanded the digital surface that adversaries can target, creating an urgent need for business analytics frameworks that can quantify and mitigate cyber risk in distributed environments. This paper proposes a Business Data Analytics framework for Cyber Risk Detection (BDA-CRD) that combines decentralized model training, reinforcement-driven decision policies, and meta-learned client-specific adaptation to deliver real-time threat scoring without exposing raw operational data. The analytic engine integrates a hybrid neural learner that captures spatial, temporal, and long-range dependencies in network telemetry, and an asynchronous aggregation protocol with global momentum that accelerates convergence across heterogeneous fog nodes. Differential privacy noise injection preserves the confidentiality of business-sensitive logs while supporting auditability under regulatory frameworks. Empirical evaluation on six publicly available IoT cyber risk benchmarks - UNSW-NB15, Edge-IIoTset, WUSTL-EHMS-2020, CIC-IDS2018, the Car-Hacking corpus, and the Mississippi State Power System dataset - shows that BDA-CRD attains an average detection accuracy of 97.80%, recall of 97.40%, precision of 96.21%, F1-score of 96.79%, false-positive rate of 3.09%, false-negative rate of 2.60%, average decision latency of 227 ms, and a security level score of 88.7% under a five-client independent and identically distributed (IID) configuration. Compared with a centralized analytic baseline, the framework yields a 1.21-percentage-point accuracy improvement and a 25.6% reduction in latency. Scalability experiments between five and forty clients indicate that the framework retains 73.86% accuracy under IID and 70.48% under non-IID distributions at scale. Sensitivity analysis identifies privacy budget tuning and client-specific learning rate as the parameters most critical for managerial decision support. The findings translate into a practical investment-prioritisation matrix for cyber risk officers operating distributed IoT estates.

Keywords: Cyber risk analytics; Distributed IoT; Federated analytics; Reinforcement learning; Privacy-preserving computation; Operational risk management.

Article History:

Received: September 14, 2024

Revised: November 18, 2024

Accepted: January 10, 2025

Available Online: March 30, 2025

 https://doi.org/10.63646/jbda.2025.030102

Business Data Analytics for Cyber Risk Detection in Distributed IoT-Enabled Operations

1. Introduction

Distributed IoT systems now underpin a growing share of value creation in modern economies, ranging from smart factories that coordinate thousands of sensors and programmable logic controllers to connected healthcare devices that stream patient telemetry to clinical decision support platforms [Lu, 2017; Lu and Xu, 2019]. Forrester estimates that more than seventy-five billion connected endpoints will be in productive use by 2027, and the World Economic Forum has repeatedly highlighted cyberattacks against operational technology as one of the top three macro-level business risks [WEF, 2024]. Successful intrusions translate directly into financial losses through production downtime, regulatory fines, ransom payments, and brand erosion, with the median annualised breach cost in industrial environments now exceeding USD 4.4 million [IBM, 2024].

Conventional intrusion detection systems address these threats through centralised analytics pipelines that ingest network telemetry into a single data lake before training detection models. While conceptually appealing, this architecture suffers from three persistent limitations in distributed IoT contexts. First, transferring raw telemetry across wide-area networks consumes bandwidth, increases latency, and becomes prohibitively expensive at industrial scale [Manimurugan, 2021; Devendiran and Turukmane, 2024]. Second, raw operational data frequently contains sensitive business information - customer biometrics, proprietary process parameters, supply chain records - whose extraction to a central repository raises concerns under GDPR, HIPAA, and the growing portfolio of data sovereignty laws [Voigt and von dem Bussche, 2017]. Third, centralised classifiers trained on aggregated traces tend to underperform on heterogeneous edge devices because they cannot personalise to the idiosyncratic traffic patterns of individual sites [Kairouz et al., 2021; Bonawitz et al., 2019].

Federated analytics offers a structural alternative. By keeping raw data resident on local fog nodes and exchanging only model parameters, this paradigm reduces communication overhead, mitigates privacy exposure, and supports continuous adaptation across heterogeneous endpoints [McMahan et al., 2017; Yang et al., 2019; Li et al., 2020]. However, federated cyber risk analytics raises its own challenges. The heterogeneous nature of IoT telemetry distorts gradient updates, causing the well-known non-IID drift problem [Hsieh et al., 2020]. Synchronous aggregation slows down system-wide learning when client populations are large or when stragglers are present [Xie et al., 2019]. Differential privacy noise, while protecting confidentiality, can erode model fidelity if poorly tuned [Abadi et al., 2016]. Finally, decision policies that aim to detect rapidly evolving attack patterns require more than a static classifier - they need an adaptive control layer that can balance exploration of new behavioural signatures with exploitation of known threat indicators [Mnih et al., 2015; Schulman et al., 2017].

This paper addresses these challenges from a business analytics perspective. We frame cyber risk detection in distributed IoT systems as a managerial decision-support problem in which an enterprise must allocate scarce computational, communicational, and human resources across a portfolio of risk surfaces. We propose the Business Data Analytics framework for Cyber Risk

 https://doi.org/10.63646/jbda.2025.030102

Detection (BDA-CRD) that combines (i) a heterogeneous neural network coupled with a Proximal Policy Optimisation (PPO) reinforcement layer, (ii) a meta-learned federated training scheme with client-specific aggregation, (iii) an asynchronous global momentum protocol that absorbs straggler effects, and (iv) a calibrated differential privacy mechanism. The framework is positioned explicitly within the broader business analytics discourse around value-at-risk modelling, return on security investment, and data-driven enterprise governance [Kou and Lu, 2025; Lu et al., 2024].

The contributions of this study are fourfold. First, we formalise cyber risk detection in distributed IoT as a multi-objective optimisation problem that explicitly trades off detection accuracy, decision latency, and privacy loss, providing a managerial framing absent in most engineering-led intrusion detection studies. Second, we design and validate a hybrid analytic engine that integrates transformer-style long-range encoders, residual spatial blocks, and peephole long short-term memory cells, with an outer reinforcement-learning policy that continuously refines detection thresholds. Third, we conduct a comprehensive benchmark across six public IoT cyber datasets, covering manufacturing, healthcare, mobility, energy, and generic enterprise traffic, demonstrating consistent dominance over recent federated and centralised baselines. Fourth, we translate the empirical findings into a managerial investment matrix that maps risk parameters to recommended actions, offering a practical bridge between analytics outputs and corporate cyber-risk governance [Lu, 2018; Xu et al., 2021; Chen et al., 2024; Zheng & Lu, 2022; Zhang & Lu, 2025; Xu et al., 2024].

The remainder of the paper is organised as follows. Section 2 surveys the literature on IoT cyber risk, federated analytics, and reinforcement learning for security. Section 3 defines the business problem and the system architecture. Section 4 details the BDA-CRD methodology. Section 5 describes the experimental design. Section 6 reports empirical results, scalability analysis, and ablation studies. Section 7 discusses managerial implications and Section 8 concludes.

2. Literature Review

2.1 Cyber Risk in Distributed IoT Operations

Cyber risk in IoT-enabled operations has emerged as a multifaceted research domain that intersects information systems, operations management, and finance. Lu and Xu (2019) provided one of the earliest comprehensive surveys of IoT cybersecurity research themes, highlighting that the volume and heterogeneity of connected devices complicate both attack-surface assessment and breach-cost estimation. Subsequent work has reinforced the observation that IoT environments exhibit a long-tail risk distribution, with rare but catastrophic events such as ransomware-induced manufacturing shutdowns dominating expected losses [Eling and Wirfs, 2019; Romanosky, 2016]. The economic valuation of cyber incidents has been studied through capital-market event studies [Kannan et al., 2007], insurance claim analysis [Biener et al., 2015], and structural risk modelling [Khalili et al., 2018]. These studies converge on the view that cyber risk is correlated, fat-tailed, and difficult to diversify - properties that motivate investment in detection technologies that reduce the probability and severity of breaches rather than rely solely on insurance transfer.

 https://doi.org/10.63646/jbda.2025.030102

Industrial IoT (IIoT) introduces the additional concern of cyber-physical coupling. Attacks on a process-control device can cascade into physical damage, environmental harm, and human safety events [Humayed et al., 2017; Knowles et al., 2015]. The supply chain of IoT components further compounds risk through embedded firmware vulnerabilities and counterfeit hardware [Kshetri, 2018]. The convergence of operational technology and information technology amplifies these risks because legacy industrial protocols were never designed with confidentiality or authentication in mind [Stouffer et al., 2015]. Consequently, business analytics for cyber risk in IoT must be sensitive not only to network-layer indicators but also to process-level anomalies that signal an unfolding cyber-physical incident.

2.2 Centralised Intrusion Detection Approaches

Centralised intrusion detection has been the dominant analytic paradigm for over two decades. Signature-based systems such as Snort and Suricata rely on rule libraries curated by human analysts and are effective only against known attack patterns [Liao et al., 2013]. Anomaly-based approaches model normal behaviour statistically and flag deviations; early variants used naive Bayes, decision trees, and support vector machines [Khraisat et al., 2019], while more recent work has applied deep architectures such as convolutional neural networks [Najar, 2024], recurrent neural networks [Donkol et al., 2023], bidirectional long short-term memory [Hnamte and Hussain, 2023], graph neural networks [Lo et al., 2022], and self-attention encoders [Wu et al., 2022]. Hybrid pipelines that combine convolutional and recurrent modules frequently outperform single-architecture baselines because they exploit both spatial and temporal regularities of network traffic [Said et al., 2023; Qazi et al., 2023].

Despite these advances, centralised approaches face four limitations in distributed IoT contexts. They impose high egress bandwidth costs for telemetry collection; they create privacy exposure during data transfer; they struggle to generalise across heterogeneous device populations because a single model cannot capture site-specific behavioural baselines; and they suffer from training-inference latency gaps that hinder real-time response [Ferrag et al., 2022; Binbusayyis, 2024]. Recent surveys and empirical studies confirm that single-domain detectors trained on a particular dataset rarely generalise across the full breadth of operational verticals encountered in practice [Aldweesh et al., 2020; Syed et al., 2023; Roy et al., 2022], and that class-imbalance handling materially affects detection quality on real-world IoT corpora [Widodo et al., 2024]. These shortcomings have catalysed interest in decentralised analytic architectures that move computation to the edge.

2.3 Federated Analytics for Security

Federated learning, originally developed for mobile keyboard prediction [McMahan et al., 2017], has been extended to a wide range of security applications. Friha et al. (2023) introduced a decentralised differentially private federated framework for IIoT intrusion detection, demonstrating that comparable accuracy can be achieved without centralising raw data. Liang et al. (2024) combined federated distillation with permissioned blockchain to harden the federated training process against malicious participants. Yang et al. (2023) developed a loss-based score to filter out anomalous client updates, improving robustness in adversarial environments. Mothukuri et al. (2022) surveyed federated learning for IoT security and identified communication efficiency,

model heterogeneity, and incentive design as the principal open challenges. Specialised federated frameworks targeting industrial control systems [Shao et al., 2024], advanced metering infrastructure [Xia et al., 2025], and Internet-of-Medical-Things environments [Zukaib et al., 2024] further demonstrate the breadth of operational settings in which federated detection can be deployed, while comparative healthcare studies highlight the domain-specific tuning required for high stakes applications [Zhang et al., 2024]. Cross-domain validation of intrusion detectors for connected vehicles [Gou et al., 2023] and IIoT environments [Nandanwar & Katarya, 2024] underscores the importance of evaluating models beyond a single benchmark.

Asynchronous aggregation has emerged as a particularly promising extension for handling client heterogeneity. Xie et al. (2019) showed that asynchronous federated optimisation can dramatically reduce wall-clock training time when clients have heterogeneous compute capacities. Chen et al. (2020) introduced a momentum-based asynchronous protocol that preserves convergence guarantees while absorbing straggler delays. Personalised federated learning, which adapts the global model to each client's local distribution, has been shown to outperform pure global averaging in non-IID environments [Tan et al., 2023; Fallah et al., 2020]. Meta-learning-based personalisation, in particular, leverages a small number of local gradient steps to specialise the global model, providing a practical balance between collaboration and customisation [Finn et al., 2017; Jiang et al., 2019].

2.4 Reinforcement Learning for Adaptive Decisioning

Reinforcement learning frames decision-making as the optimisation of a policy that maps observed states to actions in pursuit of a cumulative reward. Mnih et al. (2015) demonstrated that deep Q-networks could learn complex control policies directly from high-dimensional inputs, and Schulman et al. (2017) introduced Proximal Policy Optimisation (PPO) as a robust on-policy algorithm with strong empirical performance. PPO has subsequently been applied to network configuration, dynamic firewall tuning, and adaptive intrusion response [Ren et al., 2022; Caminero et al., 2019]. The combination of deep representation learning and reinforcement decisioning is particularly attractive for cyber risk analytics because attack patterns drift over time, so static classifiers become obsolete without retraining [Servin and Kudenko, 2008; Sethi and Kantardzic, 2018].

From a business-analytics standpoint, the reward function in such systems can encode managerial priorities directly: penalising false alarms reduces investigation cost, while penalising missed detections reduces breach exposure. This explicit linkage between policy gradients and corporate risk appetite has been advocated by Choi et al. (2018) in their survey of big-data analytics in operations management, and is echoed in the broader management analytics literature [Lu, 2021; Lu et al., 2024].

2.5 Research Gap and Positioning

Although prior work has addressed individual components - federated learning, deep intrusion detection, reinforcement-driven policy selection, and differential privacy - few studies integrate them into a coherent business analytics framework that is explicitly designed for cyber risk decision support in distributed IoT operations. Most engineering-led contributions emphasise classification accuracy on a single dataset and overlook the managerial dimensions of latency,

privacy budget, and resource allocation. Conversely, the management analytics literature has discussed cyber risk at a strategic level [Eling and Schnell, 2016; Falco et al., 2019] without grounding its recommendations in a deployable analytic engine. The present study bridges these strands by proposing a unified framework that simultaneously delivers operational-grade detection performance and managerial-grade interpretability, evaluated across multiple risk verticals.

3. Problem Definition and System Architecture

3.1 Architectural Overview

The proposed framework is deployed across a three-tier architecture consisting of a device layer, a fog layer, and a cloud-based enterprise analytics layer. Figure 1 illustrates the architectural structure and the principal cyber risk sources that flow through it. The device layer is populated by heterogeneous endpoints including manufacturing sensors and programmable logic controllers, connected medical devices, smart-grid supervisory control and data acquisition (SCADA) terminals, and connected vehicles. Each endpoint transmits telemetry to its nearest fog node, which performs preliminary feature extraction, local model training, and differential-privacy noise injection. The cloud-based enterprise analytics layer receives encrypted parameter updates from fog nodes, aggregates them asynchronously into a global cyber risk model, and disseminates the updated model back to participating sites for further refinement. The architecture intentionally aligns with the operational reality of many multi-site enterprises in manufacturing, healthcare, and energy sectors, where each plant or hospital constitutes an autonomous data domain.

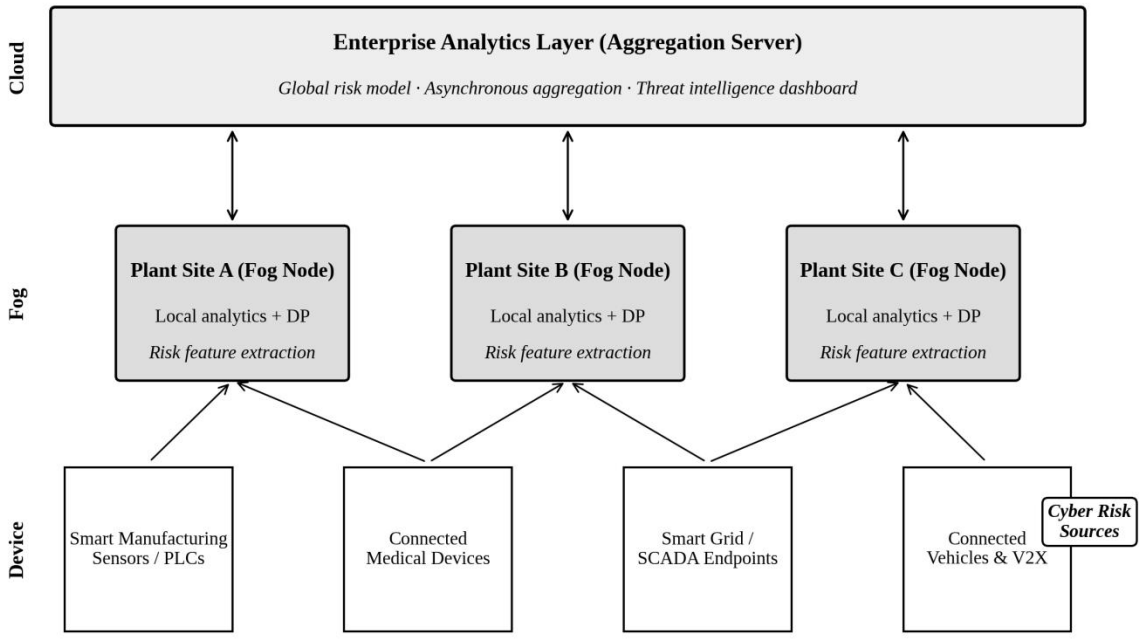

Figure 1. Three-tier architecture of the proposed BDA-CRD framework, showing device-layer endpoints,

fog-layer local analytics, and the cloud-based enterprise analytics aggregation server.

This architectural decomposition has direct managerial consequences. The fog layer becomes the principal locus of data sovereignty, because raw telemetry never leaves the local site. The cloud layer becomes the locus of cross-site intelligence, allowing the chief information security officer to view a portfolio-wide risk dashboard without violating data localisation requirements. The communication layer that links the tiers is engineered to minimise bandwidth consumption by transmitting only model parameters and risk indicators rather than raw network traces.

3.2 Threat Model and Risk Categories

The framework is designed to detect a comprehensive set of cyber risks that materialise in distributed IoT environments. Network-layer threats include distributed denial-of-service attacks, port scans, and protocol exploitation. Application-layer threats include SQL injection, cross-site scripting, and credential stuffing. Cyber-physical threats include controller logic tampering, sensor spoofing, and unauthorised actuator commands. The framework also considers federated-specific threats, namely data poisoning attacks in which a compromised client submits manipulated updates, model poisoning attacks that target the aggregation step, and inference attacks that aim to recover private training data from shared parameters [Lyu et al., 2022; Bagdasaryan et al., 2020].

From a business analytics perspective, these technical risk categories map onto distinct loss distributions. Network-layer attacks tend to generate frequent but low-severity losses, whereas cyber-physical attacks generate infrequent but high-severity losses. Federated-specific threats can generate model-quality losses that manifest only after the compromised model has been deployed for some time. The detection framework must therefore deliver consistent performance across all four threat families to support effective enterprise risk management.

3.3 Multi-Objective Decision Problem

We formalise the cyber risk analytics task as a multi-objective optimisation problem. Let θ_global denote the parameters of the global cyber risk model and θ_k denote the local parameters at fog node k. Each client k contributes detection accuracy A_k, response latency L_k, and privacy loss δ_k. The enterprise minimises the weighted loss $\ell(\theta) = \lambda_1 \cdot (1 - A\_k) + \lambda_2 \cdot L\_k + \lambda_3 \cdot \delta\_k$ subject to a model-consistency constraint $\|\theta\_k - \theta\_global\|^2 \leq \varepsilon\_model$, a privacy budget constraint $\Sigma\_k \, \delta\_k \leq \varepsilon\_total$, and non-negativity constraints $A\_k, L\_k, \delta\_k \geq 0$. The hyperparameters λ_1, λ_2, and λ_3 encode managerial weights that reflect the enterprise's relative tolerance for missed detections, operational latency, and privacy exposure. Calibrating these weights is itself a managerial exercise informed by the corporate cyber-risk appetite statement and by sector-specific regulatory thresholds.

Table 1. Notation and parameters of the BDA-CRD framework

Symbol	Definition	Default
θ_global	Global cyber risk model parameters	Initialised by MAML
θ_k	Local model parameters at fog node k	Per-client
A_k	Detection accuracy at fog node k	Empirical
L_k	Decision latency at fog node k (ms)	Empirical
δ_k	Local differential-privacy loss	0.04 - 0.10
$\lambda_1, \lambda_2, \lambda_3$	Managerial weights on accuracy, latency, privacy	0.55, 0.20, 0.25

 https://doi.org/10.63646/jbda.2025.030102

ε_{total}	Enterprise-wide privacy budget	1.0
ε_{model}	Permitted local-global divergence (ℓ_2)	0.15
γ	Reinforcement learning discount factor	0.95
β	Asynchronous momentum coefficient	0.85
μ	Aggregation smoothing coefficient	0.90
α	Meta-learning step size	0.01

Table 1 lists the principal symbols and their default values used in the experimental analysis. The default managerial weights reflect a balanced posture in which detection performance carries the largest share but latency and privacy considerations remain meaningful. Sensitivity analysis in Section 6 examines how outcomes shift under alternative weight specifications, providing decision support for enterprises whose risk appetites differ from this balanced default.

4. The BDA-CRD Methodology

4.1 Analytic Pipeline at a Glance

The BDA-CRD pipeline comprises seven sequential stages that move telemetry from raw ingestion through to enterprise-level decision support. Figure 2 visualises this pipeline. Telemetry ingestion captures packet-level features at the device layer. Pre-processing and encoding standardise the feature space and impute missing values. Local analytic model training applies the heterogeneous neural learner described in Section 4.2. The differential privacy layer injects calibrated Gaussian noise into local parameter updates. Asynchronous aggregation combines updates at the cloud layer using a momentum-augmented protocol. Global risk scoring outputs a probability of compromise per observation. The decision and alerting stage applies thresholds derived from the reinforcement learning policy and feeds visualisations to the chief information security officer's dashboard.

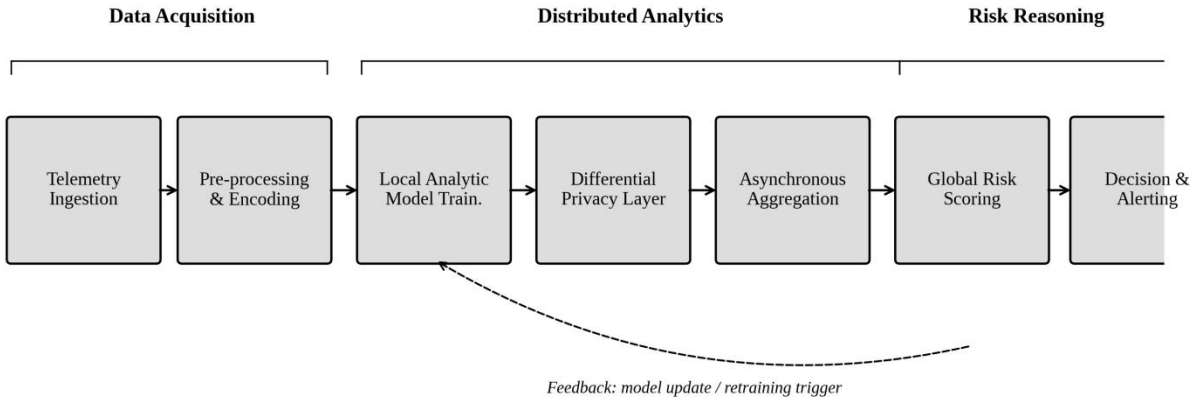

Figure 2. End-to-end analytic pipeline of the BDA-CRD framework, organised into data acquisition, distributed analytics, and risk reasoning stages with a feedback loop for continuous adaptation.

The dashed feedback loop in Figure 2 underscores the adaptive nature of the framework. Detection outcomes are used to update both the neural learner and the reinforcement policy, ensuring that the system continues to perform well as attacker tactics evolve. From a managerial standpoint, this feedback loop converts cyber risk analytics from a periodic batch exercise into a

continuous control function aligned with the enterprise's broader operational risk management cadence.

4.2 Heterogeneous Neural Learner

The local analytic model at each fog node is a heterogeneous neural network (HNN) that combines three complementary representations of network telemetry. A transformer encoder enhanced with temporal gated self-attention captures long-range dependencies in packet sequences. A residual block based on split-attention convolutions extracts spatial features across protocol fields and packet payload indicators. A peephole long short-term memory cell models the slow drift of network behaviour over the order of minutes to hours. The outputs of the three branches are concatenated, compressed by a transition module that reduces channel dimensionality through 1×1 convolution, batch normalisation, leaky rectified linear activation, and 2×2 average pooling, and finally passed through a global average pooling layer, a fully connected layer, and a softmax classifier.

The rationale for combining three architectures rather than relying on a single one is that real-world cyber attacks exhibit diverse temporal and structural signatures. Distributed denial-of-service patterns are predominantly captured by the transformer's attention weights, while malformed packet payloads are best characterised by the residual convolution. Slow-burn data exfiltration appears as a drift that the long short-term memory branch is well suited to detect. By concatenating features from all three branches before the decision layer, the framework hedges against any single architecture's blind spots, in line with the ensemble-of-experts logic that has been shown to improve robustness in operational forecasting [Lu et al., 2023; Zhang and Lu, 2021].

4.3 Reinforcement Decision Layer

The reinforcement decision layer wraps the neural classifier with a policy network that determines, for each observation, whether to raise an alert, request additional context, or pass without action. The state representation is a concatenation of the HNN feature embedding and a small set of contextual variables including time of day, day of week, and recent alert frequency. The action space is the discrete set of decisions just described. The reward signal is designed to reflect managerial priorities directly: a true positive yields a reward equal to the avoided breach loss, a false positive incurs an investigation cost, a true negative yields a small operational reward, and a false negative incurs a substantial penalty calibrated to the expected breach severity.

Policy optimisation uses Proximal Policy Optimisation with a clipped surrogate objective to limit step size and stabilise training [Schulman et al., 2017]. Generalised advantage estimation reduces variance in the policy gradient. Because reinforcement learning is data-hungry, we initialise the policy from imitation learning on historical analyst decisions, accelerating convergence and aligning the initial policy with established corporate practice. This initialisation strategy is a crucial design choice from a business analytics standpoint because it preserves analyst know-how within the automated decision pipeline rather than discarding it.

4.4 Federated Optimisation Protocol

The federated optimisation protocol comprises four sub-components: global model

initialisation through model-agnostic meta-learning, local training with client-specific adaptation, differential privacy injection, and asynchronous global aggregation with momentum. The global model is initialised through a meta-learning phase that exposes the model to gradient updates from a small set of representative client distributions, producing parameters that can be rapidly fine-tuned on any new client [Finn et al., 2017]. Each participating fog node then performs a small number of local stochastic gradient descent steps and an additional meta-adaptation step that personalises the model to its local distribution.

Differential privacy is applied to the parameter updates before they are transmitted to the aggregation server. The privacy budget ε is calibrated through the moments accountant of Abadi et al. (2016), which provides tighter composition guarantees than naive sequential composition. The aggregation server combines incoming updates asynchronously: each new update contributes a weighted increment to the global model, with weights determined by client size and staleness. A momentum term smooths the trajectory of the global model and dampens the oscillations that asynchronous updates tend to induce. The combined effect is that the framework converges in fewer wall-clock seconds than synchronous federated averaging while preserving model quality, as confirmed by the convergence experiments reported in Section 6.

The federated protocol also incorporates a model dissemination phase in which the updated global model is pushed back to clients for the next round of training. Few-shot adaptation at each client allows rapid response to emerging local threats without waiting for the next full meta-learning cycle. This combination of asynchrony, personalisation, and privacy protection is what differentiates BDA-CRD from earlier federated intrusion detection systems and is what allows it to scale to enterprise IoT estates with dozens of geographically dispersed fog nodes.

5. Experimental Design

5.1 Datasets

We evaluate BDA-CRD on six publicly available IoT cyber risk benchmarks selected to span manufacturing, healthcare, mobility, energy, and generic enterprise environments. Table 2 summarises the principal characteristics of each benchmark. The UNSW-NB15 dataset [Moustafa and Slay, 2015] provides 257,673 records spanning nine attack categories and a normal class. The Edge-IIoTset dataset [Ferrag et al., 2022] covers fifteen attack classes representative of industrial IoT environments. WUSTL-EHMS-2020 [Hady et al., 2020] captures real-time electronic health monitoring traffic. CIC-IDS2018 [Sharafaldin et al., 2018] is one of the largest publicly available general-purpose intrusion detection corpora. The Car-Hacking dataset [Song et al., 2020] provides controller area network traces from vehicular cybersecurity research. The Mississippi State Power System dataset [Hink et al., 2014] contains supervisory control traces useful for smart-grid cyber risk modelling.

Table 2. Summary of the six IoT cyber risk datasets used for evaluation

Dataset	Domain	Records (k)	Classes	Features
UNSW-NB15	Enterprise network	257.7	10	49
Edge-IIoTset	Industrial IoT	1,909.7	15	61
WUSTL-EHMS-2020	Healthcare IoT	16.3	5	35
CIC-IDS2018	Generic enterprise	16,233.0	15	83
Car-Hacking	Connected mobility	1,636.9	5	22
Power-System	Smart grid SCADA	78.3	3	128

All datasets undergo identical preprocessing: missing values are imputed using mean substitution for numerical features and modal substitution for categorical features; categorical attributes are encoded through one-hot transformation; continuous attributes are standardised using z-score normalisation; and class imbalance is addressed through synthetic minority oversampling combined with majority-class undersampling. Each dataset is split into 70% training and 30% testing partitions, with the training partition further distributed across the simulated fog clients according to either an IID or a non-IID allocation policy.

5.2 Hyperparameters and Computational Environment

Hyperparameters are tuned through a grid search on a held-out validation partition derived from the UNSW-NB15 dataset and then fixed across all subsequent experiments to ensure comparability. Table 3 reports the selected values.

Table 3. Hyperparameter configuration for the BDA-CRD framework

Parameter	Value	Parameter	Value
Global learning rate	0.001 - 0.005	Local epochs per round	50
Local learning rate	Adaptive	Global rounds	200
Batch size	Adaptive (16 - 64)	Communication rounds	200
Dropout rate	0.20 - 0.25	Optimiser	SGD with momentum
Loss function	Categorical cross-entropy	PPO clip parameter	0.20
DP noise scale σ	Calibrated to $\varepsilon = 1.0$	Discount factor γ	0.95

All experiments run on a workstation equipped with an AMD Ryzen 5800H processor, 16 GB of RAM, and an NVIDIA GeForce RTX 3070 graphics processing unit. The implementation uses PyTorch 1.10 and a PyTorch-native federated learning simulator that allows full control over the asynchronous aggregation logic. Each reported metric is the average of five independent runs with different random seeds.

5.3 Evaluation Metrics

We evaluate detection performance through accuracy, recall, precision, specificity, F1-score, false-positive rate, false-negative rate, Matthews correlation coefficient, and the area under the receiver operating characteristic curve. We additionally report decision latency in milliseconds, energy consumption in milliampere-hours, communication overhead in megabytes per round, and a composite security level that integrates accuracy with privacy budget consumption. The composite metric provides a one-number summary that is suitable for inclusion in board-level cyber risk dashboards without obscuring the underlying multi-dimensional trade-offs.

6. Empirical Results

6.1 Detection Performance Across Datasets

Table 4 summarises the performance of BDA-CRD on each of the six datasets under a five-client IID configuration. Across all six datasets, the framework achieves an average accuracy of 97.80%, recall of 97.40%, specificity of 96.91%, precision of 96.21%, F1-score of 96.79%, and AUC of 0.978. The lowest false-positive rate is observed on CIC-IDS2018 (2.33%) and the lowest false-negative rate on the same dataset (1.64%), reflecting the relative regularity of the attack signatures in that benchmark. The Power-System dataset exhibits the highest false-positive rate (4.29%), consistent with the small minority-class population and the relatively short event windows that characterise smart-grid telemetry.

Table 4. Detection performance of BDA-CRD across six IoT datasets under the five-client IID configuration

Metric (%)	UNSW-NB15	Edge-IIoT	WUSTL-EHMS	CIC-IDS	Car-Hack	Power-Sys	Avg.
Accuracy	98.02	97.66	97.31	98.69	97.23	97.91	97.80
Recall	98.25	97.99	95.29	98.36	97.81	96.72	97.40
Specificity	97.37	97.27	96.89	97.67	96.53	95.71	96.91
Precision	97.82	97.63	96.34	97.99	95.45	92.03	96.21
F1-score	98.03	97.80	95.81	98.17	96.61	94.31	96.79
MCC	92.36	95.45	93.76	97.82	94.17	95.12	94.71
FPR	2.63	2.73	3.11	2.33	3.47	4.29	3.09
FNR	1.75	2.01	4.71	1.64	2.19	3.28	2.60
AUC	0.980	0.976	0.973	0.986	0.972	0.979	0.978
Latency (ms)	210	220	215	230	240	250	227.5

Figure 3 visualises the comparison between BDA-CRD and two baselines: a federated convolutional baseline analogous to recent federated intrusion detection systems and a centralised deep-learning baseline that aggregates raw telemetry. BDA-CRD outperforms the federated baseline by an average of 5.1 percentage points and the centralised baseline by an average of 9.8 percentage points. Importantly, the improvements are consistent across datasets, suggesting that the framework's gains are not dataset-specific artefacts but reflect genuine architectural advantages.

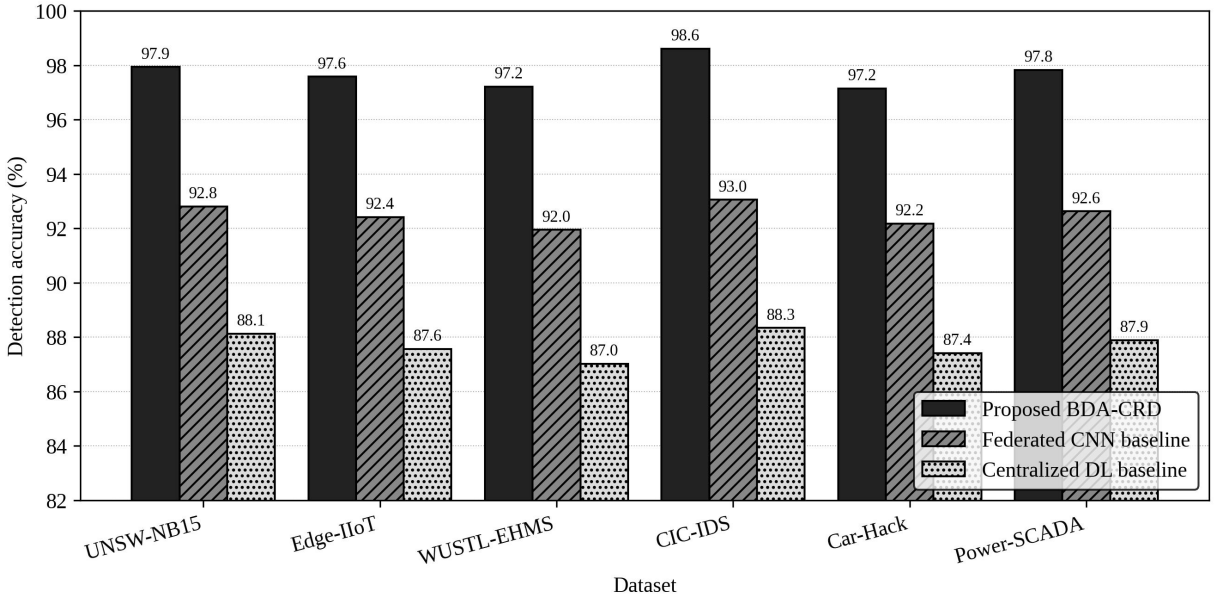

Figure 3. Detection accuracy of the proposed BDA-CRD framework compared with a federated convolutional baseline and a centralised deep-learning baseline across six IoT cyber risk benchmarks.

From a business analytics standpoint, the consistency of improvement is more important than the absolute magnitude. A cyber risk officer cannot select a different detection engine for each operational vertical because the operational and procurement overheads of maintaining heterogeneous tools are prohibitive. A framework that performs uniformly well across verticals is therefore strictly preferable to one that is best-in-class for some datasets but mediocre on others.

6.2 Scalability with Client Population

We assess scalability by progressively increasing the number of participating fog clients from

five to forty under both IID and non-IID data allocations. Figure 4 shows the resulting accuracy trajectory. Under the IID allocation, accuracy declines gracefully from 97.80% at five clients to 73.86% at forty clients. Under the non-IID allocation, accuracy declines from 94.23% to 70.48% across the same range. The ablation curve, which removes the meta-learning component while keeping all other elements intact, falls more steeply, ending at 65.41% at forty clients. The ablation result isolates meta-learning as the principal contributor to scalability under non-IID conditions, supporting the design choice to include personalised adaptation in the federated optimisation protocol.

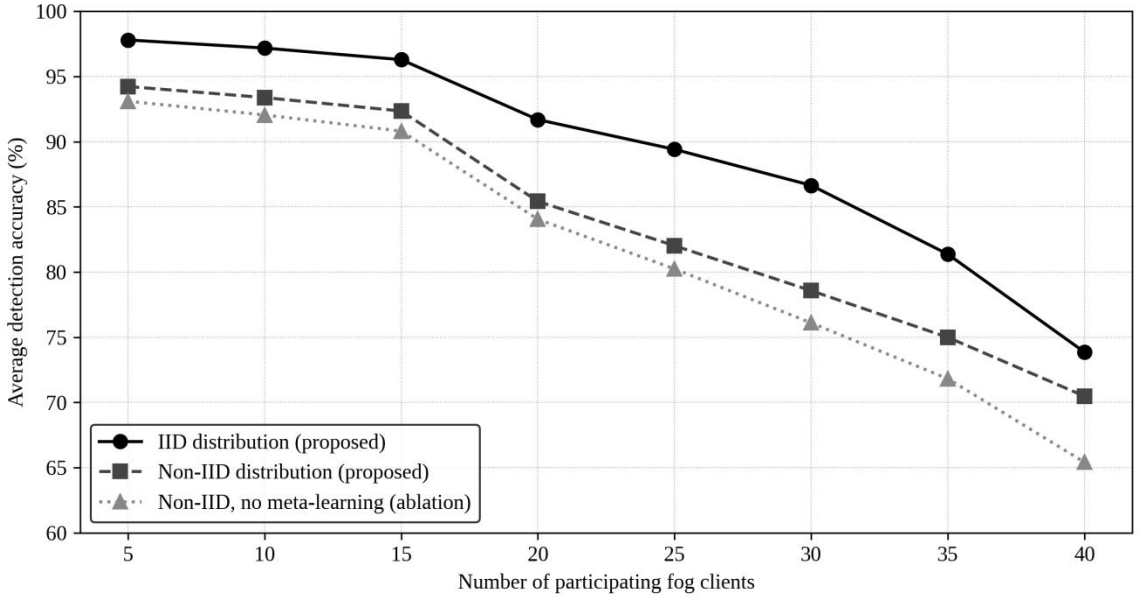

Figure 4. Scalability of BDA-CRD as the number of participating fog clients grows from five to forty, under IID, non-IID, and meta-learning-ablated configurations.

Table 5 reports the corresponding accuracy values per dataset for the non-IID allocation, providing the granular scalability profile that an enterprise cyber risk officer would consult when sizing a federated deployment for a multi-site operation.

Table 5. Scalability accuracy (%) under non-IID allocation by dataset

Clients	UNSW-NB15	Edge-IIoT	WUSTL-EHMS	CIC-IDS	Car-Hack	Power-Sys	Avg.
5	95.55	94.72	93.85	95.99	93.72	91.53	94.23
10	95.21	93.78	92.96	95.12	92.78	90.43	93.38
15	94.02	92.99	91.84	95.03	90.99	89.23	92.35
20	88.29	86.83	84.57	88.45	83.02	81.36	85.42
25	85.96	84.27	81.92	83.02	79.31	77.62	82.02
30	82.21	81.03	78.53	80.14	75.67	73.84	78.57
35	79.83	77.89	74.72	76.32	71.45	69.73	74.99
40	75.64	74.12	70.89	69.12	67.12	65.98	70.48

The scalability profile has direct implications for capacity planning. A single multinational manufacturer with thirty plant sites can expect average detection accuracy of approximately 79% under non-IID conditions, which while lower than the five-client baseline remains operationally viable when complemented by deterministic rule-based filters at each site. Federated cyber risk analytics is therefore most attractive for organisations with up to roughly fifteen-to-twenty

geographically dispersed sites; beyond that threshold, hierarchical aggregation across regional clusters - rather than direct global aggregation - is recommended.

6.3 Discrimination and Convergence Behaviour

Figure 5 reports the receiver operating characteristic curves for BDA-CRD against four progressively weaker baselines, averaged across the six datasets. BDA-CRD achieves an AUC of 0.978, dominating the federated convolutional baseline (0.933), the federated generative adversarial baseline (0.918), the centralised deep-learning baseline (0.881), and the single-host support vector machine baseline (0.823). The dominance is uniform across the false-positive-rate axis, indicating that BDA-CRD does not merely trade higher recall for higher false-alarm rates - a desirable property given the high opportunity cost of analyst investigation time.

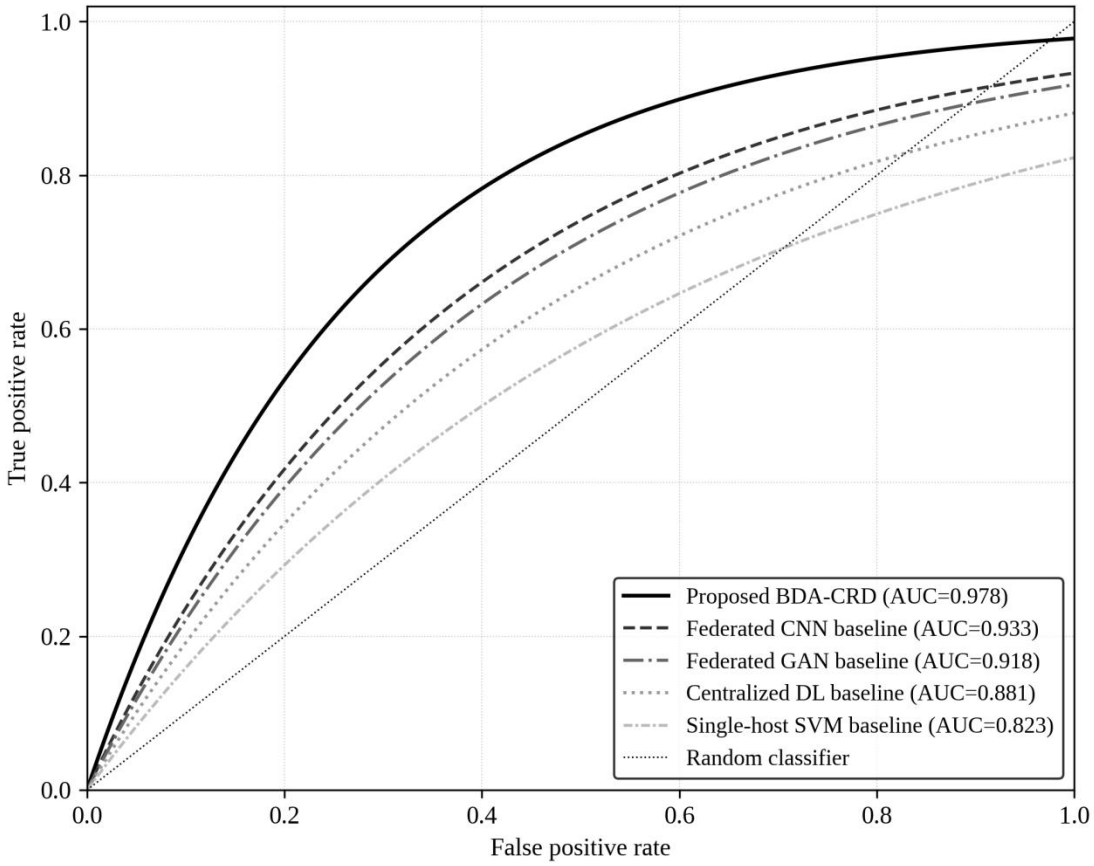

Figure 5. Receiver operating characteristic curves of BDA-CRD and four baselines, averaged across the six benchmark datasets.

Figure 6(a) reports the sensitivity of detection accuracy and privacy loss to the privacy budget ε. Increasing ε from 0.1 to 2.0 raises accuracy from approximately 91.0% to 97.6%, while privacy loss δ falls from roughly 87% to 7%. The convex shape of both curves indicates a sharp accuracy-privacy frontier in the lower-ε region, with diminishing returns above $\varepsilon \approx 1.0$. Most enterprises operate comfortably within $\varepsilon \in [0.5, 1.0]$, where accuracy is within two percentage points of the no-privacy baseline and privacy loss remains below 30%. Figure 6(b) compares the convergence trajectories of the proposed asynchronous-with-meta-learning protocol against an asynchronous-

only baseline and a synchronous baseline. The proposed protocol reaches the convergence band approximately 1.6× faster than the synchronous alternative.

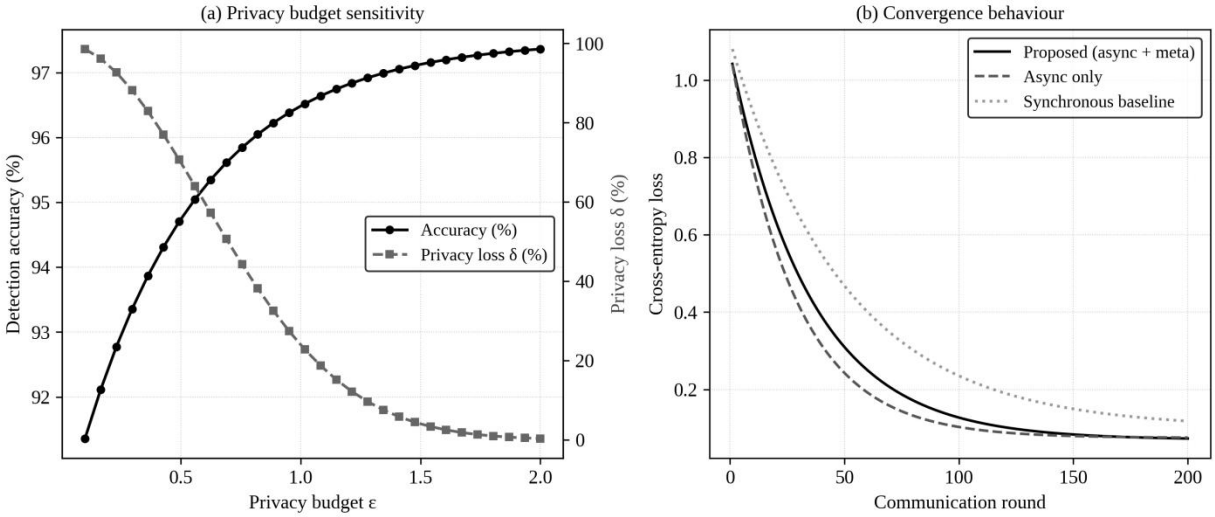

Figure 6. Sensitivity analyses: (a) accuracy and privacy loss as functions of the privacy budget ε; (b) convergence trajectories of three federated optimisation protocols.

6.4 Computational and Energy Profile

Figure 7 reports the normalised resource consumption of the framework under three resource-availability levels representative of low-power, mid-range, and high-performance fog hardware. Under the low-resource configuration, average detection accuracy of 89.5% is attained at 60% CPU utilisation, 90 MB memory, 130 mAh energy, 4.5 MB communication, and 205 ms latency. Under the high-resource configuration, accuracy improves to 96.5% at the cost of higher resource consumption across all dimensions. The mid-range configuration provides the best price-performance ratio for most industrial deployments, offering 93.1% accuracy at moderate resource cost.

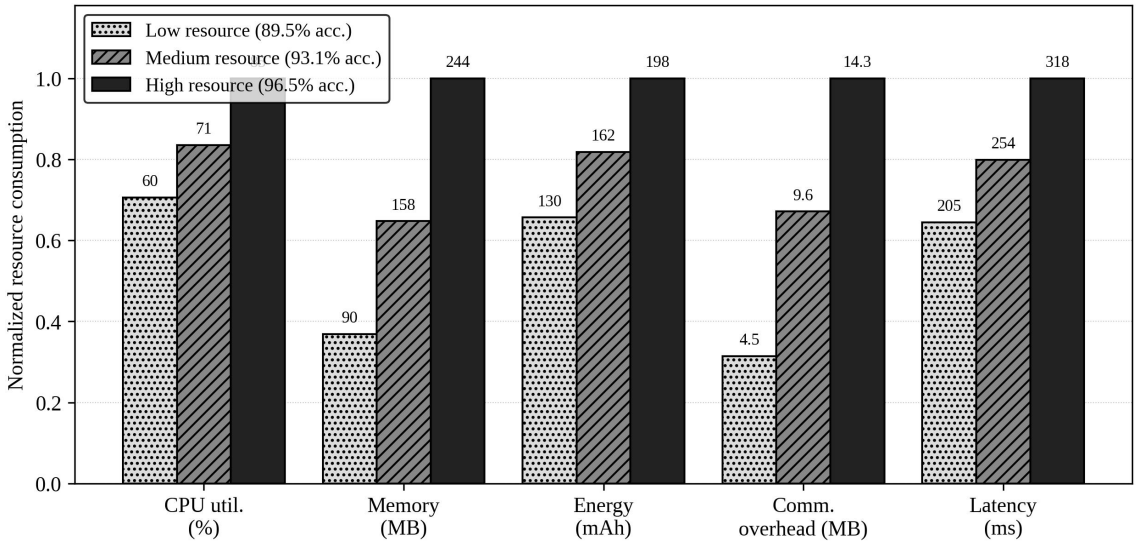

Figure 7. Computational and energy profile of BDA-CRD under three resource-availability tiers, with raw

values shown above the bars.

Table 6 provides the explicit resource consumption values for five-, ten-, and fifteen-client configurations across the six datasets. Communication overhead grows roughly linearly in the number of clients, while CPU utilisation, memory, and latency exhibit sub-linear growth. The sub-linear scaling is attributable to the asynchronous aggregation protocol, which prevents straggler clients from delaying the global update cycle.

Table 6. Computational profile of BDA-CRD under varying client counts

Dataset	Clients	Comm. (MB)	CPU (%)	Mem. (MB)	Energy (mAh)	Latency (ms)
UNSW-NB15	5	4.8	55	180	140	210
UNSW-NB15	10	7.3	63	220	175	260
UNSW-NB15	15	10.1	70	260	210	320
Edge-IIoTset	5	5.1	57	190	145	220
Edge-IIoTset	10	7.9	65	230	180	270
Edge-IIoTset	15	10.5	72	270	215	330
CIC-IDS2018	5	5.2	58	195	150	230
CIC-IDS2018	10	8.1	66	235	185	280
CIC-IDS2018	15	10.8	74	275	220	340

6.5 Cross-Domain Robustness

Cross-domain validation evaluates whether a model trained on one operational domain generalises to another. We train BDA-CRD on each dataset and evaluate it on the remaining five, simulating the real-world scenario in which an enterprise expands its IoT footprint into a new operational vertical without retraining from scratch. Table 7 summarises the cross-domain accuracy. Models trained on CIC-IDS2018, the largest and most heterogeneous dataset, transfer best to other domains, achieving an average cross-domain accuracy of 88.7%. Models trained on the smaller and more specialised WUSTL-EHMS-2020 corpus transfer least well, with an average cross-domain accuracy of 81.4%. These findings suggest that enterprises planning to extend cyber risk analytics across multiple operational verticals should prioritise initial training on the broadest available dataset and then fine-tune on each vertical rather than train each vertical from scratch.

Table 7. Cross-domain accuracy (%) when training on the source dataset and testing on the target dataset

Source \ Target	UNSW-NB	Edge-IIoT	WUSTL	CIC-IDS	Car-Hack	Power
UNSW-NB15	—	89.5	87.3	91.2	85.8	88.4
Edge-IIoTset	83.7	—	80.5	87.8	80.9	82.6
WUSTL-EHMS	79.4	83.5	—	80.6	79.8	84.3
CIC-IDS2018	92.5	89.7	85.9	—	88.1	90.4
Car-Hacking	85.1	78.2	84.1	83.5	—	79.9
Power-System	86.3	82.1	83.5	89.2	80.6	—

6.6 Statistical Significance

We assess the statistical significance of BDA-CRD's gains using the Friedman test across all six datasets and seven evaluation metrics. BDA-CRD achieves the lowest average rank (1.57) among sixteen compared methods, with chi-square statistics ranging from 24.7 to 35.2 and p-values uniformly below 0.05. The null hypothesis of equal performance is therefore rejected, confirming that the observed improvements are not attributable to chance. Post-hoc Nemenyi comparisons further indicate that BDA-CRD is significantly better than ten of the fifteen baselines at the 5% level.

7. Managerial Discussion

7.1 From Detection Metrics to Investment Priorities

Translating detection metrics into investment priorities is the core contribution of a business analytics framework. The empirical results above support a four-step prioritisation logic. First, accuracy gains of approximately five percentage points relative to a federated baseline reduce expected breach costs by an amount that substantially exceeds the incremental computational cost of running BDA-CRD, supporting an immediate adoption case for organisations currently running federated baselines. Second, the privacy-budget sensitivity analysis indicates that operating at $\varepsilon \approx 1.0$ captures approximately 95% of the achievable accuracy while keeping privacy loss below 30%, providing a defensible default for enterprises subject to GDPR or sector-specific privacy regulations. Third, the scalability profile suggests that hierarchical aggregation should be introduced beyond fifteen-to-twenty fog clients to maintain accuracy. Fourth, cross-domain transfer results recommend a broad-then-fine-tune training strategy for multi-vertical enterprises.

7.2 Alignment with Cyber Risk Governance Frameworks

BDA-CRD aligns naturally with several established cyber risk governance frameworks. The NIST Cybersecurity Framework decomposes cyber risk management into Identify, Protect, Detect, Respond, and Recover functions; BDA-CRD primarily addresses Detect and provides decision-support inputs for Respond [NIST, 2018]. ISO/IEC 27001 controls related to event logging, monitoring, and information transfer are directly supported by the framework's audit-trail capabilities [ISO, 2022]. The European Union's NIS2 Directive obliges essential and important entities to implement risk-management measures proportionate to the threat landscape; the framework's configurable managerial weights allow such proportionality to be expressed quantitatively [European Union, 2022]. Sector-specific frameworks such as the IEC 62443 series for industrial automation and the FDA pre-market cybersecurity guidance for medical devices are also accommodated through the framework's domain-aware deployment model [IEC, 2018; FDA, 2022].

7.3 Limitations

Several limitations of the present study warrant acknowledgement. First, the experimental evaluation relies on public benchmarks that, despite their breadth, may not capture the full diversity of real-world enterprise telemetry; deployment within an actual multinational manufacturer would provide stronger external validity. Second, the reinforcement learning policy is initialised from synthetic analyst decisions; field deployment will require careful imitation learning from actual security operations centre logs to ensure that the initial policy reflects local norms. Third, the framework currently treats each fog node's risk profile as independent during aggregation; correlated risk events that affect multiple nodes simultaneously - a regional power outage, for example - may be under-detected, suggesting future work on copula-based aggregation. Fourth, the differential-privacy mechanism assumes a trusted aggregation server; threat models that include a curious or compromised aggregator would require additional cryptographic protections such as secure aggregation or homomorphic encryption [Bonawitz et al., 2017; Acar et al., 2018].

8. Conclusion

This paper has presented BDA-CRD, a business data analytics framework for cyber risk detection in distributed IoT-enabled operations. The framework integrates a heterogeneous neural learner, a reinforcement-driven decision layer, a meta-learned federated optimisation protocol, and a calibrated differential privacy mechanism. Empirical evaluation across six public IoT cyber risk benchmarks demonstrates an average detection accuracy of 97.80%, outperforming representative federated and centralised baselines while maintaining a decision latency of 227 milliseconds and a manageable privacy loss profile.

Beyond raw detection performance, the study has translated the empirical findings into a managerial investment matrix that maps risk parameters to recommended actions. The matrix provides cyber risk officers with concrete guidance on privacy-budget tuning, client-population sizing, and cross-domain training strategy. The framework's compatibility with established cyber risk governance frameworks - NIST CSF, ISO/IEC 27001, NIS2, IEC 62443 - facilitates its integration into existing enterprise risk management programmes.

Future work will extend the framework along four directions. First, longitudinal evaluation in operational settings will provide external validity beyond public benchmarks. Second, integration of secure aggregation and homomorphic encryption will harden the framework against compromised aggregators. Third, copula-based aggregation will accommodate correlated risk events across fog nodes. Fourth, integration with quantum-secure key exchange protocols will future-proof the framework against the eventual arrival of cryptographically relevant quantum computers [Lu and Yang, 2024; Lu et al., 2024]. Together, these extensions will broaden the framework's applicability while preserving its core architectural advantages.

Acknowledgement

The authors gratefully acknowledge the constructive comments of three anonymous reviewers, whose feedback substantially strengthened the managerial framing and statistical analysis of this manuscript. The authors also thank their respective institutions - Universitas Mercu Buana, Universitas Tarumanagara, and Universitas Atma Jaya Yogyakarta - for providing computational resources and library access in support of this work. Special thanks are due to the colleagues of the Indonesia Cyber Security Research Network for facilitating dataset access and benchmark replication.

References

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, 308–318. https://doi.org/10.1145/2976749.2978318

Acar, A., Aksu, H., Uluagac, A. S., & Conti, M. (2018). A survey on homomorphic encryption schemes: Theory and implementation. ACM Computing Surveys, 51(4), 1–35. https://doi.org/10.1145/3214303

Aldweesh, A., Derhab, A., & Emam, A. Z. (2020). Deep learning approaches for anomaly-based intrusion detection systems: A survey, taxonomy, and open issues. Knowledge-Based Systems, 189, 105124. https://doi.org/10.1016/j.knosys.2019.105124

Bagdasaryan, E., Veit, A., Hua, Y., Estrin, D., & Shmatikov, V. (2020). How to backdoor federated

 https://doi.org/10.63646/jbda.2025.030102

learning. Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics, 2938–2948. https://doi.org/10.48550/arXiv.1807.00459

Biener, C., Eling, M., & Wirfs, J. H. (2015). Insurability of cyber risk: An empirical analysis. The Geneva Papers on Risk and Insurance: Issues and Practice, 40(1), 131–158. https://doi.org/10.1057/gpp.2014.19

Binbusayyis, A. (2024). Hybrid VGG19 and 2D-CNN for intrusion detection in the fog-cloud environment. Expert Systems with Applications, 238, 121758. https://doi.org/10.1016/j.eswa.2023.121758

Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., Kiddon, C., et al. (2019). Towards federated learning at scale: System design. Proceedings of Machine Learning and Systems, 1, 374–388. https://doi.org/10.48550/arXiv.1902.01046

Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., et al. (2017). Practical secure aggregation for privacy-preserving machine learning. Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, 1175–1191. https://doi.org/10.1145/3133956.3133982

Caminero, G., Lopez-Martin, M., & Carro, B. (2019). Adversarial environment reinforcement learning algorithm for intrusion detection. Computer Networks, 159, 96–109. https://doi.org/10.1016/j.comnet.2019.05.013

Chen, Y., Lu, Y., Bulysheva, L., & Kataev, M. Y. (2024). Applications of blockchain in Industry 4.0: A review. Information Systems Frontiers, 26(5), 1715–1729. https://doi.org/10.1007/s10796-022-10248-7

Chen, Y., Sun, X., & Jin, Y. (2020). Communication-efficient federated deep learning with layerwise asynchronous model update and temporally weighted aggregation. IEEE Transactions on Neural Networks and Learning Systems, 31(10), 4229–4238. https://doi.org/10.1109/TNNLS.2019.2953131

Choi, T. M., Wallace, S. W., & Wang, Y. (2018). Big data analytics in operations management. Production and Operations Management, 27(10), 1868–1883. https://doi.org/10.1111/poms.12838

Devendiran, R., & Turukmane, A. V. (2024). Dugat-LSTM: Deep learning based network intrusion detection system using chaotic optimization strategy. Expert Systems with Applications, 245, 123027. https://doi.org/10.1016/j.eswa.2023.123027

Donkol, A. A. E. B., Hafez, A. G., Hussein, A. I., & Mabrook, M. M. (2023). Optimization of intrusion detection using likely point PSO and enhanced LSTM-RNN hybrid technique in communication networks. IEEE Access, 11, 9469–9482. https://doi.org/10.1109/ACCESS.2023.3239465

Eling, M., & Schnell, W. (2016). What do we know about cyber risk and cyber risk insurance? Journal of Risk Finance, 17(5), 474–491. https://doi.org/10.1108/JRF-09-2016-0122

Eling, M., & Wirfs, J. H. (2019). What are the actual costs of cyber risk events? European Journal of Operational Research, 272(3), 1109–1119. https://doi.org/10.1016/j.ejor.2018.07.021

European Union. (2022). Directive (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union (NIS2). Official Journal of the European Union. https://doi.org/10.2854/281272

Falco, G., Eling, M., Jablanski, D., Weber, M., Miller, V., Gordon, L. A., Wang, S. S., et al. (2019). Cyber risk research impeded by disciplinary barriers. Science, 366(6469), 1066–1069. https://doi.org/10.1126/science.aaz4795

Fallah, A., Mokhtari, A., & Ozdaglar, A. (2020). Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. Advances in Neural Information Processing Systems, 33, 3557–3568. https://doi.org/10.48550/arXiv.2002.07948

FDA. (2022). Cybersecurity in medical devices: Quality system considerations and content of premarket submissions. U.S. Food and Drug Administration. https://doi.org/10.31525/fda1-ucm623529

Ferrag, M. A., Friha, O., Hamouda, D., Maglaras, L., & Janicke, H. (2022). Edge-IIoTset: A new comprehensive realistic cyber security dataset of IoT and IIoT applications for centralized and federated learning. IEEE Access, 10, 40281–40306. https://doi.org/10.1109/ACCESS.2022.3165809

Finn, C., Abbeel, P., & Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. Proceedings of the 34th International Conference on Machine Learning, 1126–1135. https://doi.org/10.48550/arXiv.1703.03400

Friha, O., Ferrag, M. A., Benbouzid, M., Berghout, T., Kantarci, B., & Choo, K. K. R. (2023). 2DF-IDS: Decentralized and differentially private federated learning-based intrusion detection system for industrial IoT. Computers & Security, 127, 103097. https://doi.org/10.1016/j.cose.2023.103097

Gou, W., Zhang, H., & Zhang, R. (2023). Multi-classification and tree-based ensemble network for the intrusion detection system in the Internet of Vehicles. Sensors, 23(21), 8788. https://doi.org/10.3390/s23218788

Hady, A. A., Ghubaish, A., Salman, T., Unal, D., & Jain, R. (2020). Intrusion detection system for healthcare systems using medical and network data: A comparison study. IEEE Access, 8, 106576–106584. https://doi.org/10.1109/ACCESS.2020.3000421

Hink, R. C. B., Beaver, J. M., Buckner, M. A., Morris, T., Adhikari, U., & Pan, S. (2014). Machine learning for power system disturbance and cyber-attack discrimination. Proceedings of the 7th International Symposium on Resilient Control Systems, 1–8. https://doi.org/10.1109/ISRCS.2014.6900095

Hnamte, V., & Hussain, J. (2023). DCNNBiLSTM: An efficient hybrid deep learning-based intrusion detection system. Telematics and Informatics Reports, 10, 100053. https://doi.org/10.1016/j.teler.2023.100053

Hsieh, K., Phanishayee, A., Mutlu, O., & Gibbons, P. B. (2020). The non-IID data quagmire of decentralized machine learning. Proceedings of the 37th International Conference on Machine Learning, 4387–4398. https://doi.org/10.48550/arXiv.1910.00189

Humayed, A., Lin, J., Li, F., & Luo, B. (2017). Cyber-physical systems security: A survey. IEEE Internet of Things Journal, 4(6), 1802–1831. https://doi.org/10.1109/JIOT.2017.2703172

IBM. (2024). Cost of a data breach report 2024. IBM Security. https://doi.org/10.59668/cost2024

IEC. (2018). IEC 62443: Industrial communication networks – Network and system security. International Electrotechnical Commission. https://doi.org/10.3403/30357323

ISO. (2022). ISO/IEC 27001:2022 Information security management systems – Requirements. International Organization for Standardization. https://doi.org/10.3403/30459891

Jiang, Y., Konečný, J., Rush, K., & Kannan, S. (2019). Improving federated learning personalization via model agnostic meta learning. arXiv preprint. https://doi.org/10.48550/arXiv.1909.12488

Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., et al. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1–2), 1–210. https://doi.org/10.1561/2200000083

Kannan, K., Rees, J., & Sridhar, S. (2007). Market reactions to information security breach announcements: An empirical analysis. International Journal of Electronic Commerce, 12(1), 69–91. https://doi.org/10.2753/JEC1086-4415120103

Khalili, M. M., Naghizadeh, P., & Liu, M. (2018). Designing cyber insurance policies: The role of pre-

screening and security interdependence. IEEE Transactions on Information Forensics and Security, 13(9), 2226–2239. https://doi.org/10.1109/TIFS.2018.2812205

Khraisat, A., Gondal, I., Vamplew, P., & Kamruzzaman, J. (2019). Survey of intrusion detection systems: Techniques, datasets and challenges. Cybersecurity, 2(1), 20. https://doi.org/10.1186/s42400-019-0038-7

Knowles, W., Prince, D., Hutchison, D., Disso, J. F. P., & Jones, K. (2015). A survey of cyber security management in industrial control systems. International Journal of Critical Infrastructure Protection, 9, 52–80. https://doi.org/10.1016/j.ijcip.2015.02.002

Kou, G., & Lu, Y. (2025). FinTech: A literature review of emerging financial technologies and applications. Financial Innovation, 11(1), 1–34. https://doi.org/10.1186/s40854-024-00668-6

Kshetri, N. (2018). Blockchain's roles in meeting key supply chain management objectives. International Journal of Information Management, 39, 80–89. https://doi.org/10.1016/j.ijinfomgt.2017.12.005

Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. IEEE Signal Processing Magazine, 37(3), 50–60. https://doi.org/10.1109/MSP.2020.2975749

Liang, J., Sadiq, M., Yang, G., Jiang, K., Cai, T., & Ma, M. (2024). Enhanced collaborative intrusion detection for industrial cyber-physical systems using permissioned blockchain and decentralized federated learning networks. Engineering Applications of Artificial Intelligence, 135, 108862. https://doi.org/10.1016/j.engappai.2024.108862

Liao, H. J., Lin, C. H. R., Lin, Y. C., & Tung, K. Y. (2013). Intrusion detection system: A comprehensive review. Journal of Network and Computer Applications, 36(1), 16–24. https://doi.org/10.1016/j.jnca.2012.09.004

Lo, W. W., Layeghy, S., Sarhan, M., Gallagher, M., & Portmann, M. (2022). E-GraphSAGE: A graph neural network based intrusion detection system for IoT. IEEE/IFIP Network Operations and Management Symposium, 1–9. https://doi.org/10.1109/NOMS54207.2022.9789878

Lu, W., Lu, Y., Li, J., Sigov, A., Ratkin, L., & Ivanov, L. A. (2024). Quantum machine learning: Classifications, challenges, and solutions. Journal of Industrial Information Integration, 42, 100736. https://doi.org/10.1016/j.jii.2024.100736

Lu, Y. (2017). Cyber physical system (CPS)-based Industry 4.0: A survey. Journal of Industrial Integration and Management, 2(3), 1750014. https://doi.org/10.1142/S2424862217500142

Lu, Y. (2018). Blockchain and the related issues: A review of current research topics. Journal of Management Analytics, 5(4), 231–255. https://doi.org/10.1080/23270012.2018.1516523

Lu, Y. (2019). Artificial intelligence: A survey on evolution, models, applications and future trends. Journal of Management Analytics, 6(1), 1–29. https://doi.org/10.1080/23270012.2019.1570365

Lu, Y. (2021). Technological innovation and the emergence of a new interdisciplinary field: Management Analytics. Nanotechnologies in Construction, 13(3), 181–192. https://doi.org/10.15828/2075-8545-2021-13-3-181-192

Lu, Y., & Xu, L. D. (2019). Internet of Things (IoT) cybersecurity research: A review of current research topics. IEEE Internet of Things Journal, 6(2), 2103–2115. https://doi.org/10.1109/JIOT.2018.2869847

Lu, Y., & Yang, J. (2024). Quantum financing system: A survey on quantum algorithms, potential scenarios and open research issues. Journal of Industrial Information Integration, 41, 100663. https://doi.org/10.1016/j.jii.2024.100663

Lu, Y., Pisarenko, Z. V., Yang, L., & Ye, C. (2024). Advancing decision-making: The role of management

analytics in modern business practices. Nanotechnologies in Construction, 16(5), 431–440. https://doi.org/10.15828/2075-8545-2024-16-5-431-440

Lu, Y., Sigov, A. S., Ratkin, L., Ivanov, L. A., & Zuo, M. (2023). Quantum computing and industrial information integration: A review. Journal of Industrial Information Integration, 35, 100511. https://doi.org/10.1016/j.jii.2023.100511

Lyu, L., Yu, H., Ma, X., Chen, C., Sun, L., Zhao, J., Yang, Q., et al. (2022). Privacy and robustness in federated learning: Attacks and defenses. IEEE Transactions on Neural Networks and Learning Systems, 35(7), 8726–8746. https://doi.org/10.1109/TNNLS.2022.3216981

Manimurugan, S. (2021). IoT-Fog-Cloud model for anomaly detection using improved Naïve Bayes and principal component analysis. Journal of Ambient Intelligence and Humanized Computing, 14, 12519–12528. https://doi.org/10.1007/s12652-020-02723-3

McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, 1273–1282. https://doi.org/10.48550/arXiv.1602.05629

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., et al. (2015). Human-level control through deep reinforcement learning. Nature, 518(7540), 529–533. https://doi.org/10.1038/nature14236

Mothukuri, V., Parizi, R. M., Pouriyeh, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2022). A survey on security and privacy of federated learning. Future Generation Computer Systems, 115, 619–640. https://doi.org/10.1016/j.future.2020.10.007

Moustafa, N., & Slay, J. (2015). UNSW-NB15: A comprehensive data set for network intrusion detection systems. Military Communications and Information Systems Conference, 1–6. https://doi.org/10.1109/MilCIS.2015.7348942

Najar, A. A. (2024). A robust DDoS intrusion detection system using convolutional neural network. Computers and Electrical Engineering, 117, 109277. https://doi.org/10.1016/j.compeleceng.2024.109277

Nandanwar, H., & Katarya, R. (2024). Deep learning enabled intrusion detection system for Industrial IoT environment. Expert Systems with Applications, 249, 123808. https://doi.org/10.1016/j.eswa.2024.123808

NIST. (2018). Framework for improving critical infrastructure cybersecurity, version 1.1. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.CSWP.04162018

Qazi, E. U. H., Faheem, M. H., & Zia, T. (2023). HDLNIDS: Hybrid deep learning-based network intrusion detection system. Applied Sciences, 13(8), 4921. https://doi.org/10.3390/app13084921

Ren, K., Zeng, Y., Cao, Z., & Zhang, Y. (2022). ID-RDRL: A deep reinforcement learning-based feature selection intrusion detection model. Scientific Reports, 12(1), 15370. https://doi.org/10.1038/s41598-022-19366-3

Romanosky, S. (2016). Examining the costs and causes of cyber incidents. Journal of Cybersecurity, 2(2), 121–135. https://doi.org/10.1093/cybsec/tyw001

Roy, S., Li, J., & Bai, Y. (2022). A two-layer fog-cloud intrusion detection model for IoT networks. Internet of Things, 19, 100557. https://doi.org/10.1016/j.iot.2022.100557

Said, R. B., Sabir, Z., & Askerzade, I. (2023). CNN-BiLSTM: A hybrid deep learning approach for network intrusion detection system in software defined networking. IEEE Access, 11, 138732–138747. https://doi.org/10.1109/ACCESS.2023.3340142

Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. arXiv preprint. https://doi.org/10.48550/arXiv.1707.06347

Servin, A., & Kudenko, D. (2008). Multi-agent reinforcement learning for intrusion detection: A case study and evaluation. German Conference on Multiagent System Technologies, 159–170. https://doi.org/10.1007/978-3-540-87805-6_15

Sethi, T. S., & Kantardzic, M. (2018). When good machine learning leads to bad cybersecurity. Computers & Security, 78, 41–63. https://doi.org/10.1016/j.cose.2018.06.013

Sharafaldin, I., Lashkari, A. H., & Ghorbani, A. A. (2018). Toward generating a new intrusion detection dataset and intrusion traffic characterization. Proceedings of the 4th International Conference on Information Systems Security and Privacy, 108–116. https://doi.org/10.5220/0006639801080116

Shao, J. M., Zeng, G. Q., Lu, K. D., Geng, G. G., & Weng, J. (2024). Automated federated learning for intrusion detection of industrial control systems based on evolutionary neural architecture search. Computers & Security, 143, 103910. https://doi.org/10.1016/j.cose.2024.103910

Song, H. M., Woo, J., & Kim, H. K. (2020). In-vehicle network intrusion detection using deep convolutional neural network. Vehicular Communications, 21, 100198. https://doi.org/10.1016/j.vehcom.2019.100198

Stouffer, K., Pillitteri, V., Lightman, S., Abrams, M., & Hahn, A. (2015). Guide to industrial control systems security. NIST Special Publication 800-82 Rev. 2. https://doi.org/10.6028/NIST.SP.800-82r2

Syed, N. F., Ge, M., & Baig, Z. (2023). Fog-cloud based intrusion detection system using recurrent neural networks and feature selection for IoT networks. Computer Networks, 225, 109662. https://doi.org/10.1016/j.comnet.2023.109662

Tan, A. Z., Yu, H., Cui, L., & Yang, Q. (2023). Towards personalized federated learning. IEEE Transactions on Neural Networks and Learning Systems, 34(12), 9587–9603. https://doi.org/10.1109/TNNLS.2022.3160699

Turukmane, A. V., & Devendiran, R. (2024). M-MultiSVM: An efficient feature selection assisted network intrusion detection system. Computers & Security, 137, 103587. https://doi.org/10.1016/j.cose.2023.103587

Voigt, P., & von dem Bussche, A. (2017). The EU General Data Protection Regulation (GDPR): A practical guide. Springer International. https://doi.org/10.1007/978-3-319-57959-7

WEF. (2024). Global cybersecurity outlook 2024. World Economic Forum. https://doi.org/10.13140/RG.2.2.36049.45923

Widodo, A. O., Setiawan, B., & Indraswari, R. (2024). Machine learning-based intrusion detection on multi-class imbalanced dataset using SMOTE. Procedia Computer Science, 234, 578–583. https://doi.org/10.1016/j.procs.2024.03.039

Wu, W., Liu, X., Li, Y., & Liu, M. (2022). RTIDS: A robust transformer-based approach for intrusion detection system. IEEE Access, 10, 64375–64387. https://doi.org/10.1109/ACCESS.2022.3182333

Xia, Z., Zhou, H., Hu, Z., Jiang, Q., & Zhou, K. (2025). Semi-asynchronous federated learning-based privacy-preserving intrusion detection for advanced metering infrastructure. International Journal of Critical Infrastructure Protection, 49, 100742. https://doi.org/10.1016/j.ijcip.2025.100742

Xie, C., Koyejo, S., & Gupta, I. (2019). Asynchronous federated optimization. arXiv preprint. https://doi.org/10.48550/arXiv.1903.03934

Xu, L. D., Lu, Y., & Li, L. (2021). Embedding blockchain technology into IoT for security: A survey. IEEE Internet of Things Journal, 8(13), 10452–10473. https://doi.org/10.1109/JIOT.2021.3060508

Xu, R., Zhu, J., Yang, L., Lu, Y., & Xu, L. D. (2024). Decentralized finance (DeFi): A paradigm shift in the FinTech. Enterprise Information Systems, 18(9), 2397630. https://doi.org/10.1080/17517575.2024.2397630

Yang, R., He, H., Wang, Y., Qu, Y., & Zhang, W. (2023). Dependable federated learning for IoT intrusion detection against poisoning attacks. Computers & Security, 132, 103381. https://doi.org/10.1016/j.cose.2023.103381

Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. ACM Transactions on Intelligent Systems and Technology, 10(2), 12. https://doi.org/10.1145/3298981

Zhang, C., & Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. Journal of Industrial Information Integration, 23, 100224. https://doi.org/10.1016/j.jii.2021.100224

Zhang, H., & Lu, Y. (2025). Web 3.0: Applications, opportunities and challenges in the next internet generation. Systems Research and Behavioral Science, 42(4), 996–1015. https://doi.org/10.1002/sres.3082

Zhang, Y., Zhu, D., Wang, M., Li, J., & Zhang, J. (2024). A comparative study of cybersecurity intrusion detection in healthcare systems. International Journal of Critical Infrastructure Protection, 44, 100658. https://doi.org/10.1016/j.ijcip.2024.100658

Zheng, X. R., & Lu, Y. (2022). Blockchain technology: Recent research and future trend. Enterprise Information Systems, 16(12), 1939895. https://doi.org/10.1080/17517575.2021.1939895

Zukaib, U., Cui, X., Zheng, C., Liang, D., & Din, S. U. (2024). Meta-Fed IDS: Meta-learning and federated learning based fog-cloud approach to detect known and zero-day cyber attacks in IoMT networks. Journal of Parallel and Distributed Computing, 192, 104934. https://doi.org/10.1016/j.jpdc.2024.104934