International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(20s) 2026 pp. 1672-1693
https://doi.org/10.70917/ijcisim-2026-6088
Article

Intelligent Cybersecurity Analytics for Supply Chain Resilience: Integrating AI, Machine Learning, and Big Data for Proactive Threat Hunting

Muhammad Noman Amjad[1], Muhammad Akram[2], Afrasiyab Ali[3], Asma Javaid[4]

[1]Department of Information Technology, Washington University of Science and Technology (WUST), United States. Email: noman5802@gmail.com
[2]Harper Community College, 1200 Algonquin Rd, Palatine, IL 60067, United States
Email: m.akarm.achakzai@gmail.com
[3]Computer Science, Business, Stellar Career College, United States. Email: afrasiyabali2003@gmail.com
[4]Department of Software Engineering, The University of Azad Jammu and Kashmir, Muzaffarabad, Pakistan. Email: asma.javaid@uajk.edu.pk

Abstract: Today, supply chains are intricate digital ecosystems that include suppliers, logistics providers, industrial Internet of Things (IIoT) assets, cloud platforms and software, and adversaries are leveraging these trust relationships as a first step to gaining access. Signature-driven and purely supervised detectors provide excellent protection against known tactics but suffer dramatically when an attacker introduces tactics not seen in historical data – which is exactly what proactive threat hunting must account for. This paper introduces ICA-SCR, an Intelligent Cybersecurity Analytics framework for Supply Chain Resilience, which integrates a big data ingestion and feature layer, a supervised gradient-boosting channel for known threats, a benign-calibrated unsupervised anomaly channel (autoencoder and isolation forest) for unknown threats, and a Dual-Channel Hunting Fusion (DCHF) rule that transforms both channels into analyst-ready hunt priorities, along with a SHAP-based explainability layer. The framework is tested on two publicly available datasets, NSL-KDD (148,517 records) and UNSW-NB15 (257,673 records), against seven supervised and two unsupervised baselines under six protocols, which include binary detection, cross-validation, novel-attack hunting, leave-one-category-out zero-day simulation, multi-class attribution and streaming scalability. On NSL-KDD, where the test partition includes 17 attack types not present in the training partition, ICA-SCR improved the F1-score from 0.761 to 0.838 and the MCC from 0.615 to 0.706 compared to its own supervised channel, while the false-positive rate increased from 2.81% to 3.14% (McNemar $\chi^2 = 1323.9$, $p < 0.001$). The number of attacks of new types detected more than doubled (30.3% to 64.5%). The supervised channel already generalised well on UNSW-NB15, where attack categories share strong traffic signatures, and fusion added a small but statistically significant false-alarm cost, which delineates when the anomaly channel adds value. The scoring pipeline achieved approximately 43,000 records per second on a 2-core CPU. These findings reveal that supervised and benign-calibrated anomaly analytics is a feasible and comprehensible path to proactive threat hunting in supply chain settings and quantify the compromises practitioners must make.

Keywords: Supply Chain Resilience; Cybersecurity Analytics; Proactive Threat Hunting; Machine Learning; Anomaly Detection; Big Data; Explainable AI; Zero-Day Detection

1. Introduction

Global supply chains are now highly interconnected cyberphysical systems. Procurement, warehousing, transportation and production planning are now reliant on common enterprise resource planning (ERP) platforms, warehouse and transport management systems, supplier portals, IIoT sensors and a host of third party software. While

Copyright: © 2026 by the authors

this connectivity enhances efficiency, it also makes every partner an attack surface: sensitive financial and operational information flows constantly between organisations [1] and end-to-end visibility in such networks is challenging with a siloed monitoring approach [2]. Thus, effective defence will require a mix of machine analytics and human expertise in the form of collaborative intelligence arrangements [3].

These ecosystems are under attack and the attacks are multi-stage and temporal. They span a long dwell time, with reconnaissance, credential theft, and lateral movement and exfiltration occurring in tandem with legitimate partner traffic [4]. Cyber threat intelligence (CTI) mining can uncover indicators of such campaigns before they reach maturity [5] and artificial intelligence (AI) is now considered as a fundamental brick of resilient digital ecosystems [6]. The urgency is backed up by industry evidence. According to the 2025 Verizon Data Breach Investigations Report, which analyzed over 22,000 incidents and confirmed 12,195 breaches, third-party involvement has doubled to 30% of breaches and exploitation of vulnerabilities has increased by 34% [28]. According to IBM's 2025 Cost of a Data Breach study, the first vector in approximately 15% of breaches was via a supply chain compromise, and supply chain and vendor compromises took the longest of any breach type to identify and contain, averaging 267 days. The same study revealed that organisations that extensively adopted security AI and automation reduced the breach lifecycle by 80 days and saved approximately US$1.9 million per breach [29]. Attacks on software supply chains in 2025 and 2026, including the compromise of popular open-source packages to harvest CI/CD tokens and cloud secrets, demonstrate that attackers are targeting the trust relationships that supply chains depend on..

Threat hunting is the proactive, hypothesis-driven process of looking for adversaries that have slipped past automated controls. It is evolving along structured, unstructured and situational lines with the help of AI, especially when it comes to advanced persistent threats (APTs) [7] and autonomous hunting agents have been proposed as a future paradigm [8]. Basic explanations of AI in cybersecurity emphasize the importance of the simultaneous development of automation and intelligent decision-making and explainability [9]. Holistic architectures that combine AI, machine learning (ML) and data analytics for detection and response [10] and predictive big-data analytics with behavioural ML for IIoT [11] reinforces the same conclusion: the value is in the combination of these capabilities and not any given model.

However, three gaps still exist. First, most of the ML-based detectors are trained on the same attack vocabulary, and therefore the accuracy is not indicative of the ability to discover novel tactics, techniques and procedures (TTPs) which is the main goal of hunting. Second, security studies for supply chains are either conceptual in nature, suggesting architectures, or report accuracy of a single model, with few quantifying the effect of a combined supervised and unsupervised design on detection of unseen attacks, false-alarm load and analyst workload. Third, the operational characteristics, which are the most critical factors in whether analytics can be resilient in practice, such as explainability for the analysts, calibrated alert budgets and streaming throughput, are rarely measured in combination.

This paper addresses these gaps with the following contributions:

- **A layered framework (ICA-SCR)** which integrates supply chain telemetry, big-data feature engineering, a known threat supervised channel, a benign calibrated anomaly channel, explainable prioritisation of hunts and a resilience feedback loop.
- **A Dual-Channel Hunting Fusion (DCHF) rule** that transforms the scores of autoencoder and isolation-forest into benign-traffic percentiles. This provides an interpretable and budgeted false-alarm rate for the anomaly channel that can be fused with supervised probabilities without requiring retraining.
- **A multi-protocol empirical evaluation** on two public benchmarks using seven supervised and two unsupervised baselines for novel-attack hunting, leave-one-category-out zero-day simulation, multi-class attribution, statistical significance testing, SHAP explainability and streaming scalability.
- **An evidence-based account of when fusion helps and when it does not.** The two sets of data provide conflicting results and the paper converts these to advice for the supply chain security operations.

The rest of this paper is structured as follows. In Section 2, related works are reviewed. The ICA-SCR framework is presented in Section 3. The experimental design is described in Section 4. The results are reported in Section 5, implications and limitations are discussed in Section 6 and a conclusion is drawn in Section 7.

2. Related Work

2.1 AI and Machine Learning for Cybersecurity Analytics

To coordinate intelligent hunting and response across distributed infrastructure, multi-agent, cloud-native defence frameworks have been proposed [12]. Documents cite the evolution from static signatures to learned behavioural models in surveys of AI and ML for next generation of detection [13] and predictive analytics is increasingly presented as the foundation of anticipatory rather than reactive defence [14]. In IoT environment, proactive hunting is related to the sustainability of the ecosystem itself [15]. Platform-oriented approaches integrate deep learning, transfer learning and behavioural analytics to enhance enterprise resilience against adversaries using AI technologies [16]. A recent survey of deep, hybrid and resilience-oriented architectures, which combines complementary learners, has found that, in operational scenarios, the hybrid architectures are the most promising, though they are rarely tested against unknown threats [17].

2.2 Supply Chain Cyber Resilience and Big Data

In the supply chain area, AI and blockchain are fused to provide secure sustainable supply chains with tamper evident records and intelligent verification [18]. As the volume of data and number of partners increase, big-data platforms must be complemented by forward-looking cyber defence for supply chain optimisation [19]. Collective-intelligence mechanisms enable several stakeholders to collectively manage cybersecurity risks throughout supply ecosystems [20]. Deep learning has been used to secure the cyber-physical systems of Supply Chain 4.0 [21]. Cloud-native enterprise data platforms, such as [22] and AI-powered threat intelligence for real-time anomaly detection of cloud infrastructure [23] focus on the hosting layer on which most supply chain applications are now hosted.

Tan et al. [24] give a detailed analysis of APTs using supply chain vulnerabilities and list the detection of stealthy and previously unknown behaviour as a key open challenge. Experimental studies of predictive network intelligence have shown that ML can be used for proactive detection, but also report sensitivity to training distribution [25]. Logistics resilience strategies highlight the need for multiple layers of control, ongoing monitoring and coordination with partners [26] and design science research demonstrates that AI-enabled predictiveness can be directly integrated into operational processes to enhance the resilience of the supply chain [27].

2.3 Research Gap

Table 1 provides an overview of representative studies across the dimensions addressed in this work. While it is well accepted in the literature that AI, ML and big data are key to supply chain cyber resilience, few studies consider both supervised and unsupervised channels in one calibrated hunting pipeline and examine explicitly for unseen attack types, while providing explainability and throughput in addition to accuracy. ICA-SCR is created to fill this gap.

TABLE 1 Positioning of ICA-SCR relative to representative literature

Study	Primary focus	Supply chain context	Unseen-threat evaluation	XAI	Scalability reported
[2], [3], [20]	Collaborative / collective intelligence for monitoring and risk	Yes	No	Partial	No
[4]	HMM-based temporal threat	Yes	Partial	No	No

	recognition and automated hunting				
[5], [7], [8]	CTI mining; AI-driven and autonomous threat hunting	Indirect	Discussed	Partial	No
[11], [15], [21]	Big-data / deep learning for IIoT and CPS security	Yes (IIoT, SC 4.0)	No	No	Partial
[18], [19], [26]	Blockchain, big data and strategy for SC security	Yes	No	No	No
[24]	Survey of supply chain APTs	Yes	Open challenge	–	–
[12], [22], [23]	Cloud-native AI defence and threat intelligence	Indirect	No	Partial	Partial
ICA-SCR (this work)	Calibrated dual-channel analytics for proactive hunting	Yes	Yes (novel types + LOCO)	Yes (SHAP)	Yes

Note: The characterisation reflects each study's stated scope. LOCO = leave-one-category-out; XAI = explainable AI.

3. Proposed ICA-SCR Framework

3.1 Architecture Overview

The six layers of ICA-SCR are depicted in Fig. 1. L1 gathers telemetry data from the supply chain attack surface, such as supplier network links and third-party logistics (3PL) network links, ERP, WMS and TMS logs, IIoT and OT gateways, cloud and SaaS APIs and software build pipelines. L2 is a large scale ingestion and feature layer, which aggregates events into flow-level records in micro-batches, and encodes categorical attributes and normalises heavy tailed counters. L3a is a supervised channel that identifies known TTPs, while L3b is an unsupervised anomaly channel that is only trained with benign behaviour and can detect deviations from normal even if there is no labelled example available. L4 merges both channels in alerts and hunt priorities. L5 adds SHAP explanations and builds a ranked hunt queue, and L6 completes the picture: analyst-validated outcomes are used to retrain the supervised channel, refreshed benign baselines are used to recalibrate the anomaly threshold, and CTI, software bill of materials (SBOM) data and vendor risk scores are used to enrich hypotheses. This organisation mirrors the human–machine collaboration model proposed in [2] [3] and the hunting paradigm of [7], [8].

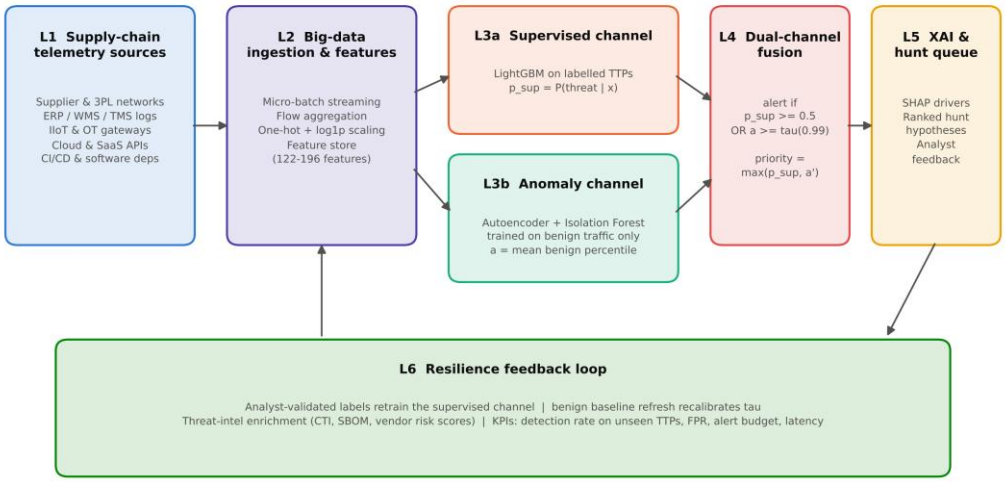


Fig. 1. Layered architecture of the proposed ICA-SCR framework. Supply chain telemetry (L1) is converted into features (L2) and scored by a supervised channel for known TTPs (L3a) and a benign-calibrated anomaly channel for unknown TTPs (L3b). Dual-Channel Hunting Fusion (L4) produces alerts and priorities, which are explained and queued for analysts (L5) and fed back into model maintenance (L6).

3.2 Problem Formulation

Let *x_ ∈ ℝ^d^ be a feature vector describing one network flow or session observed at a supply chain boundary, with label _y* ∈ {0, 1} (benign or malicious). The training set *D_ = {(_x*$_i$, *y_~i~)} covers a set of attack types _A*$_{train}$. At deployment, traffic may contain attack types *A_~test~* ⊄ *_A*$_{train}$. Proactive hunting asks for a scoring function *s_(_x)* and decision rule that (i) maximise detection on *A_~test~* \ *_A*$_{train}$, (ii) keep the false-positive rate (FPR) on benign traffic within an analyst budget, and (iii) rank alerts so that limited analyst time is spent on the most probable threats.

3.3 Big-Data Ingestion and Feature Engineering

Records are processed in micro-batches to contain memory and latency. Categorical protocol, service and connection-state attributes are one-hot encoded over the union of observed values, resulting in 122 and 196 features for NSL-KDD and UNSW-NB15, respectively. The heavy-tailed distributions of traffic counters (bytes, packets, durations, loads, etc.) destabilise both distance-based and reconstruction-based models, and every non-negative numeric feature v is transformed as

$$v' = \ln(1 + \max(v, 0)) \tag{1}$$

The linear and neural baselines additionally receive z-score standardisation fitted on the training partition. The anomaly channel uses min–max scaling fitted on benign data and clipped to [0, 1] at inference, so that unseen extremes saturate rather than distort reconstruction.

3.4 Supervised Channel (L3a)

The supervised channel is a LightGBM gradient-boosted decision tree ensemble, which has been used for its accuracy on tabular data, as well as its histogram-based training speed and native support for exact tree SHAP explanations [30], [31]. After *M* boosting rounds the model outputs

$$p_{sup}(x) = \sigma\left(\Sigma_{m=1\ldots M}\, \eta \cdot f_m(x) \right), \quad \sigma(z) = 1 / (1 + e^{-z}) \tag{2}$$

where f_m is the *m_-th regression tree and _η* the learning rate. The channel is trained with binary log-loss on all labelled training records.

3.5 Benign-Calibrated Anomaly Channel (L3b)

The anomaly channel is only trained on benign records, thus the concept of normality is not based on any attack label. Two complementary detectors are used which are operated simultaneously. The reconstruction error of an under-complete autoencoder g_ with a 64–16–64 ReLU architecture that is trained to reconstruct scaled benign inputs _z provides the anomaly evidence.

$$e(x) = (1/d) \, \| \, z - g(z) \, \|^2 \tag{3}$$

An isolation forest [32] with 200 trees provides a second, partition-based view, whose anomaly evidence *r_(_x*) is the negated average path-length score. The two raw scores have incomparable units, so each is mapped to its empirical cumulative distribution function (ECDF) on a held-out 20% benign reference set *B*:

$$F_e(x) = |\{ \, b \in B : e(b) \le e(x) \, \}| \, / \, |B|, \quad a(x) = \tfrac{1}{2} \, [\, F_e(x) + F_r(x) \,] \tag{4}$$

The calibrated score *a_(_x*) ∈ [0, 1] reads as the fraction of known-benign traffic that looks more normal than *x*. This interpretation is useful for analysts and allows the alerting threshold to be stated as a false-alarm budget.

3.6 Dual-Channel Hunting Fusion (L4)

DCHF raises an alert when either channel has sufficient evidence:

$$\hat{y}(x) = \mathbb{1}[\, p_{sup}(x) \ge 0.5 \ \vee \ a(x) \ge \tau \,], \quad \tau = Q_q(\, a(B) \,) \tag{5}$$

where Q_q is the *q_-quantile of calibrated scores on the benign reference set. With _q* = 0.99 the anomaly channel contributes, by construction, about 1% additional false alarms on in-distribution benign traffic. The learned thresholds were $\tau = 0.974$ for NSL-KDD and $\tau = 0.952$ for UNSW-NB15. For hunt queue ordering, each record receives the priority

$$\pi(x) = \max\{ \, p_{sup}(x), \ \tfrac{1}{2} \cdot \mathbb{1}[a(x) \ge \tau] + \tfrac{1}{2} \cdot (a(x) - \tau)/(1 - \tau) \cdot \mathbb{1}[a(x) \ge \tau] \, \} \tag{6}$$

so that confident supervised detections and strong anomalies compete on a common [0, 1] scale. The OR-fusion is deliberately simple. It needs no labels for unseen attacks, adds no trainable parameters that could overfit to known TTPs, and can be audited by analysts. Algorithm 1 summarises the procedure.

Algorithm 1. ICA-SCR training and Dual-Channel Hunting Fusion

Step	**Operation**
Input	Labelled training flows $D = \{(x_i, y_i)\}$; benign quantile q (default 0.99)
1	Encode categorical fields (one-hot) and apply $v' = \ln(1+v)$ to numeric counters (Eq. 1)
2	Train supervised channel p_{sup} (LightGBM) on all of D (Eq. 2)
3	Split benign records $D_0 = \{x_i : y_i = 0\}$ into fit (80%) and reference B (20%) sets
4	Fit min–max scaler, autoencoder g and isolation forest on the fit set (Eq. 3)
5	Build ECDFs F_e, F_r on B; compute calibrated score $a(\cdot)$ (Eq. 4)
6	Set $\tau \leftarrow$ q-quantile of a(B)
7	For each incoming micro-batch: compute p_{sup}, a; alert $\hat{y}$ (Eq. 5); priority π (Eq. 6)
8	Attach SHAP explanation of p_{sup} for each alert; push to ranked hunt queue
9	Periodically retrain p_{sup} with analyst-validated labels; refresh B and recalibrate τ (L6)
Output	Alerts $\hat{y}$, hunt priorities π, and explanations for analysts

3.7 Explainability Layer (L5)

Each alert from the supervised channel is explained with TreeSHAP [30], which decomposes the log-odds output into additive per-feature contributions φ_j satisfying

$$\text{logit} \, p_{sup}(x) = \varphi_0 + \Sigma_{j=1 \ldots d} \, \varphi_j(x) \tag{7}$$

Global importance is defined as the mean |φj| computed for a sample of test records. Anomaly-channel alerts are provided with their benign percentile a_(_x) which provides a direct statement about the unusualness of the flow relative to normal operation. These two combined provide a why and a how unusual for each hunt lead, serving the explainability requirement emphasized in [9] and [17].

4. Experimental Design

4.1 Datasets and Supply Chain Relevance

As mentioned in [24] there is no public data set that contains labelled attacks across an operational multi-tier supply chain. Two popular network intrusion benchmarks were thus chosen: one that represents the network border between partners where intrusions into the supply chain can be detected, and one that has official splits that enable to evaluate threats that cannot be seen.

TABLE 2 Characteristics of the evaluation datasets (official train/test partitions)

Property	**NSL-KDD [33]**	**UNSW-NB15 [34]**
Training records	125,973	175,341
Test records	22,544	82,332
Raw / encoded features	41 / 122	42 / 196
Attack share (train / test)	46.5% / 56.9%	68.1% / 55.1%
Attack categories	DoS, Probe, R2L, U2R	Analysis, Backdoor, DoS, Exploits, Fuzzers, Generic, Reconnaissance, Shellcode, Worms
Attack types unseen in training	17 types, 3,750 test records	None (simulated with LOCO protocol)
Supply chain analogue	Partner-facing legacy services; credential and remote-access abuse (R2L/U2R)	Modern hybrid traffic; exploits, backdoors and worms typical of software supply chain campaigns

NSL-KDD is a cleaned version of KDD Cup 99 which doesn't contain duplicate records [33]. It has 3,750 records from 17 attack types that never appear in the training set, including mscan, apache2, processtable and httptunnel, making it a natural testbed for searching for new TTPs. UNSW-NB15 [34] consists of synthetic traffic created using IXIA PerfectStorm tool which combines realistic modern benign traffic with nine modern attack families. The class distributions shown in Fig. 2 are highly skewed for both datasets. The minority categories (U2R, Worms and Shellcode) represent the less common but high impact intrusions that supply chain hunters are most interested in discovering.

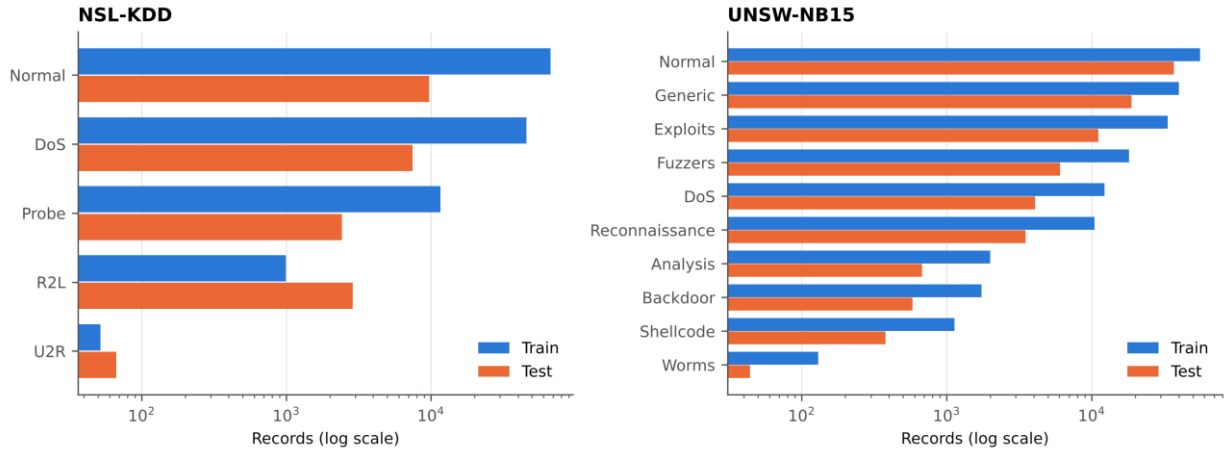


Fig. 2. Class distribution of the training and test partitions (log scale). Both benchmarks are imbalanced, and the NSL-KDD test partition has a markedly larger share of R2L and U2R records than training, many of them from unseen attack types.

4.2 Evaluation Protocols

Six protocols were designed, each mapped to a resilience question:

- **E1 – Binary detection:** train on the official training partition and test on the official test partition, with the addition of data from the partner network (can the framework separate threats from benign partner traffic?).
- **E2 – Cross-validated robustness:** stratified 5-fold cross-validation on training partition for the three best supervised learners (what is the stability of the in-distribution performance, and what is the difference to the shifted test distribution?).
- **E3 – Novel-attack hunting:** detection rate on the 17 attack types (NSL-KDD) not included in training (Does the framework detect new TTPs?).
- **E4 – Leave-one-category-out (LOCO) zero-day simulation:** Train the supervised channel without one of the nine UNSW-NB15 attack categories and test on it (How robust is detection to an all new attack family?
- **E5 – Multi-class attribution and explainability:** ten class classification and Shapley analysis (can analysts tell what is happening and why it was flagged) for UNSW-NB15.
- **E6 – Hunting efficiency and scalability:** alert composition, unseen-threat capture under an alert budget, and streaming throughput and latency for micro-batches (is the pipeline operationally viable?).

4.3 Baselines and Hyperparameters

ICA-SCR was compared to seven supervised learners: logistic regression (LR), decision tree (DT), random forest (RF), extra trees (ET), XGBoost, LightGBM and a multilayer perceptron (MLP) and the two unsupervised detectors of the anomaly channel separately. The LightGBM baseline separates the effect of fusion since LightGBM is the supervised channel in ICA-SCR. Hyperparameters were set a priori to values that are commonly used and are not tuned on the test partitions (see Table 3). Random seed 42 was used for all stochastic components.

TABLE 3 Model configurations

Model	**Key hyperparameters**
LR	L2 penalty, C = 1.0, max_iter = 300, standardised inputs

DT	max_depth = 20
RF / ET	150 trees, unrestricted depth
XGBoost	300 rounds, max_depth = 8, η = 0.1, subsample = 0.9, colsample = 0.8, hist
LightGBM (L3a)	400 rounds, 63 leaves, η = 0.05, bagging 0.9, feature fraction 0.8
MLP	128–64 ReLU, Adam, early stopping, ≤ 40 epochs, standardised inputs
Autoencoder (L3b)	64–16–64 ReLU, Adam, early stopping, ≤ 60 epochs, ≤ 40,000 benign records
Isolation Forest (L3b)	200 trees, contamination = auto, fitted on benign records
DCHF (L4)	Benign quantile q = 0.99; supervised threshold 0.5; 20% benign reference set

4.4 Evaluation Metrics

TP, TN, FP and FN are true and false positives and negatives, respectively, and the following measures are reported: accuracy, precision, recall or detection rate (DR), F1-score, area under the ROC curve (AUC), false-positive rate and Matthews correlation coefficient (MCC) which is resistant to class imbalance:

$$F1 = 2 \cdot P \cdot R / (P + R), \quad FPR = FP / (FP + TN) \qquad (8)$$

$$MCC = (TP \cdot TN - FP \cdot FN) / \sqrt{[(TP+FP)(TP+FN)(TN+FP)(TN+FN)]} \qquad (9)$$

Paired differences between ICA-SCR and its supervised channel on the same test records were assessed with McNemar's test with continuity correction, $\chi^2 = (|b - c| - 1)^2 / (b + c)$, where *b_ and _c* count records that only one of the two models classifies correctly.

4.5 Implementation Environment

All experiments were conducted using Python 3, scikit-learn 1.9, LightGBM 4.7, XGBoost 3.4 and SHAP 0.53 on a cloud container with 2 virtual CPU cores, 7 GB of RAM, and no GPU. This relatively simple setup was intentionally selected to represent the resources within the reach of small and medium-sized supply chain partners. The full code, raw outputs and figure scripts are available as supplementary material.

5. Results

5.1 Binary Threat Detection (E1)

TABLE 4 Binary detection performance on the NSL-KDD test partition (22,544 records)

Model	**Acc.**	**Prec.**	**Rec.**	**F1**	**AUC**	**FPR**	**MCC**	**Train (s)**	**Infer. (µs/rec)**
LR	0.7482	0.9143	0.6153	0.7356	0.7901	0.0762	0.5491	1.2	0.2
DT	0.7932	0.9660	0.6599	0.7842	0.8147	0.0307	0.6392	1.3	0.3
RF	0.7669	0.9669	0.6115	0.7492	0.9607	0.0277	0.6023	5.8	5.8
ET	0.7888	0.9692	0.6497	0.7779	0.9573	0.0273	0.6344	6.5	7.5
XGBoost	0.7932	0.9682	0.6582	0.7837	0.9715	0.0285	0.6402	5.8	3.4
LightGBM	0.7757	0.9672	0.6273	0.7610	0.9679	0.0281	0.6148	4.0	16.3
MLP	0.7971	0.9259	0.6995	0.7970	0.9167	0.0739	0.6257	10.7	0.7
Autoencoder	0.7586	0.9266	0.6255	0.7468	0.9502	0.0655	0.5701	32.7	–
IsolationForest	0.7445	0.9844	0.5600	0.7139	0.9522	0.0117	0.5802	32.7	–

ICA-SCR (DCHF)	**0.8373**	**0.9688**	**0.7379**	**0.8377**	**0.9687**	**0.0314**	**0.7059**	**36.8**	**25.9**

Note: Unsupervised models use the calibrated score with threshold 0.99. ICA-SCR training time includes the

LightGBM channel. Inference times are wall-clock times on two CPU cores; the unsupervised rows are not timed separately because they are part of the ICA-SCR pipeline.

Table 4 indicates that for NSL-KDD all supervised models have high precision (0.914–0.969) while having low recall (0.612–0.700). They lack the ability to capture the training attack vocabulary and fail to capture a significant portion of the test attack vocabulary, many of which are new. ICA-SCR improves recall from 0.627 (ICA-SCR's LightGBM channel) to 0.738, while maintaining precision at 0.969. Its F1 score is 0.838 and MCC score is 0.706, which are the highest among all 10 models, 4.1 points better than the best baseline (MLP, F1 = 0.797) and 7.7 points better than LightGBM. The cost is not high: the FPR increases from 2.81% to 3.14%, which is significantly lower than the 7.4–7.6% of the MLP and LR baselines. The improvement is not due to chance as confirmed by the McNemar's test. Of the records on which the two models disagree, ICA-SCR is correct on 1,419 and LightGBM on only 32 ($\chi^2 = 1323.9$, $p < 0.001$).

TABLE 5 Binary detection performance on the UNSW-NB15 test partition (82,332 records)

Model	**Acc.**	**Prec.**	**Rec.**	**F1**	**AUC**	**FPR**	**MCC**	**Train (s)**	**Infer. (µs/rec)**
LR	0.8387	0.7833	0.9776	0.8697	0.9671	0.3314	0.6933	7.6	0.2
DT	0.8674	0.8273	0.9595	0.8885	0.8971	0.2455	0.7393	2.8	0.2
RF	0.8701	0.8158	0.9870	0.8933	0.9794	0.2730	0.7531	15.9	7.1
ET	0.8671	0.8132	0.9849	0.8909	0.9738	0.2772	0.7469	15.0	11.2
XGBoost	0.8710	0.8191	0.9829	0.8935	0.9833	0.2661	0.7532	12.5	3.3
LightGBM	0.8738	0.8215	0.9847	0.8957	0.9847	0.2622	0.7587	5.8	12.0
MLP	0.8659	0.8163	0.9762	0.8891	0.9792	0.2691	0.7416	39.7	1.7
Autoencoder	0.7918	0.9737	0.6391	0.7717	0.8971	0.0211	0.6399	54.5	–
IsolationForest	0.5562	0.9483	0.2052	0.3374	0.8085	0.0137	0.2940	54.5	–
ICA-SCR (DCHF)	**0.8650**	**0.8107**	**0.9848**	**0.8893**	**0.9840**	**0.2818**	**0.7432**	**60.7**	**24.4**

Note: Columns as in Table 4.

A different picture emerges from UNSW-NB15 (Table 5). Tree ensembles reach AUCs of 0.974–0.985 and recall above 0.98, and LightGBM achieves the best F1 (0.896) and MCC (0.759). ICA-SCR matches its recall (0.985) and AUC (0.984), but its F1 is 0.6 points lower because FPR rises from 26.2% to 28.2%. In this instance the McNemar comparison has a preference for the supervised channel (725 vs. 7 discordant records, $\chi^2 = 702.3$, $p < 0.001$). Thus the anomaly channel introduces false alarms, and fails to detect attacks missed by the supervised channel. The relatively high FPR of all the models on UNSW-NB15 is a known consequence of the distribution shift between the official partitions of the dataset (Section 5.2) and not a peculiarity of ICA-SCR. The ROC curves and the F1 and MCC of all models are shown in Fig. 3 and Fig. 4.

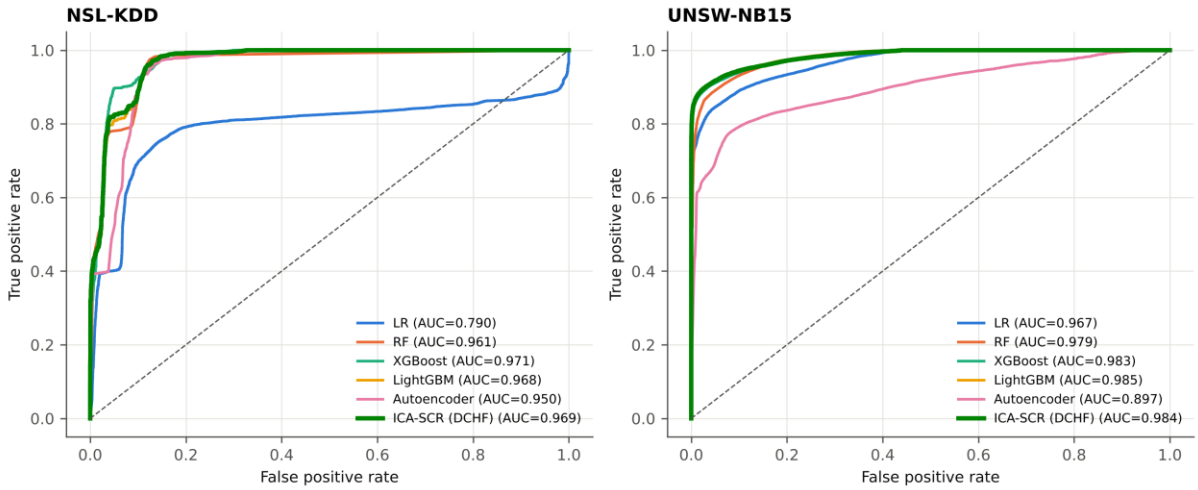


Fig. 3. ROC curves of representative models on NSL-KDD (left) and UNSW-NB15 (right). ICA-SCR (thick green) tracks the strongest ensembles on both datasets. On NSL-KDD it dominates the low-FPR region that matters for analyst workload, apart from XGBoost between FPR ≈ 0.03 and 0.10.

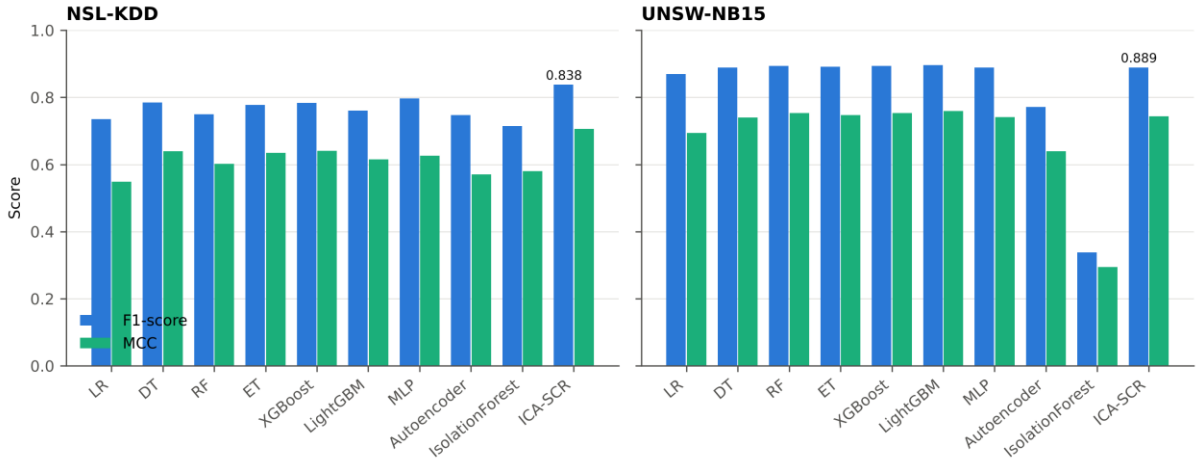


Fig. 4. F1-score and Matthews correlation coefficient for all ten models. ICA-SCR gives the best F1 and MCC on NSL-KDD, where unseen attacks dominate the misses. On UNSW-NB15 it is within 0.6 F1 points of the best supervised model.

5.2 Cross-Validated Robustness and Distribution Shift (E2)

TABLE 6 Stratified 5-fold cross-validation on the training partitions (mean ± standard deviation)

Model	Accuracy	Precision	Recall	F1	AUC	FPR
NSL-KDD · RF	0.9989 ± 0.0002	0.9994 ± 0.0002	0.9982 ± 0.0004	0.9988 ± 0.0002	1.0000 ± 0.0000	0.0005 ± 0.0002
NSL-KDD · XGBoost	0.9991 ± 0.0002	0.9993 ± 0.0002	0.9987 ± 0.0004	0.9990 ± 0.0002	1.0000 ± 0.0000	0.0006 ± 0.0001
NSL-KDD · LightGBM	0.9993 ± 0.0002	0.9995 ± 0.0001	0.9990 ± 0.0003	0.9993 ± 0.0002	1.0000 ± 0.0000	0.0004 ± 0.0001

UNSW-NB15 · RF	0.9597 ± 0.0013	0.9618 ± 0.0010	0.9796 ± 0.0012	0.9706 ± 0.0009	0.9938 ± 0.0003	0.0828 ± 0.0022
UNSW-NB15 · XGBoost	0.9621 ± 0.0015	0.9655 ± 0.0011	0.9792 ± 0.0013	0.9723 ± 0.0011	0.9946 ± 0.0003	0.0745 ± 0.0023
UNSW-NB15 · LightGBM	0.9620 ± 0.0013	0.9653 ± 0.0011	0.9795 ± 0.0008	0.9723 ± 0.0009	0.9946 ± 0.0003	0.0751 ± 0.0024

The in-distribution performance is very high and stable as shown in Table 6, where the F1 score of all three ensembles on NSL-KDD is greater than 0.998 with at most 0.0002 standard deviation, and the F1 score on UNSW-NB15 is approximately 0.972 ± 0.001. These are in contrast to the official test figures (F1 of 0.75-0.78 and 0.89-0.90 respectively). The difference between the performance on NSL-KDD and UNSW-NB15 is > 20 F1 points and ~ 8 points, respectively, which indicates the accuracy gap between the deployment traffic and the training traffic. The most significant result of the robustness analysis for supply chain operators is that cross-validated accuracy, which is the main focus of the literature, can greatly overestimate field performance when supply chain partners, software, or adversary TTPs vary.

5.3 Proactive Hunting of Unseen Attack Types (E3)

The NSL-KDD test partition enables one to distinguish the attacks of types present in the training set from the attacks of types that are entirely new. Fig. 5(a) indicates that fusion does not affect the detection of known types at all significantly (76.1% to 77.6%), but increases the detection of unseen types by more than two times (30.3% to 64.5%). Fig. 5(b) shows the gain by category. Probe detection rises from 74.0% to 98.1%, DoS from 81.6% to 89.5%, U2R from 17.9% to 43.3% and R2L from 5.6% to 13.5%.

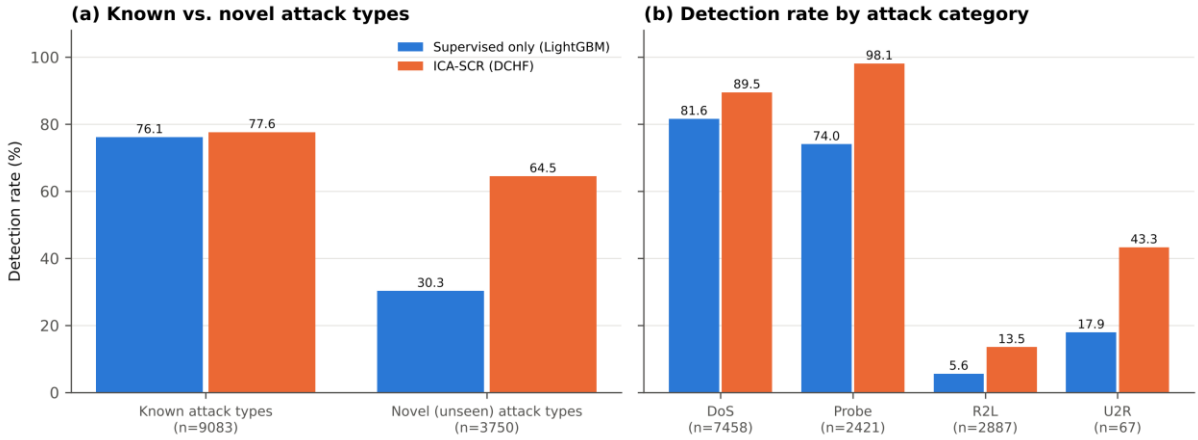


Fig. 5. Detection rates on NSL-KDD for the supervised channel alone and for ICA-SCR: (a) known versus unseen attack types; (b) by attack category. The anomaly channel more than doubles the detection of unseen attack types.

TABLE 7 Detection of individual attack types that never appear in the NSL-KDD training data

Unseen attack type	Test records	Supervised DR (%)	ICA-SCR DR (%)	Δ (pp)
mscan	996	37.1	95.7	+58.5
apache2	737	40.3	88.7	+48.4
processtable	685	20.3	52.7	+32.4

snmpguess	331	0.0	0.0	+0.0
saint	319	99.1	99.1	+0.0
mailbomb	293	0.0	0.0	+0.0
snmpgetattack	178	0.0	0.0	+0.0
httptunnel	133	2.3	85.0	+82.7
named	17	5.9	29.4	+23.5
ps	15	20.0	33.3	+13.3
sendmail	14	0.0	0.0	+0.0
xterm	13	15.4	46.2	+30.8
xlock	9	0.0	0.0	+0.0
xsnoop	4	25.0	50.0	+25.0
sqlattack	2	50.0	50.0	+0.0
udpstorm	2	100.0	100.0	+0.0
worm	2	0.0	0.0	+0.0

The gains per attack type are shown in Table 7 and Fig. 6. The biggest gains come from the resource-exhaustion attacks on apache2 (40.3% to 88.7%) and processtable (20.3% to 52.7%), and the multi-host scanning tool mscan (37.1% to 95.7%), along with the covert channel httptunnel (2.3% to 85.0%). Covert tunnelling and wide area scanning are indeed the hallmarks of a compromised supplier being used as a beachhead by an adversary [7] [24]. There are still some content-level attacks that both channels, snmpguess, snmpgetattack, mailbomb and sendmail are unable to detect. They do not have any flow characteristics that distinguish them from benign SNMP or SMTP sessions, and thus would require payload or protocol-semantic features to detect them (Section 6.3).

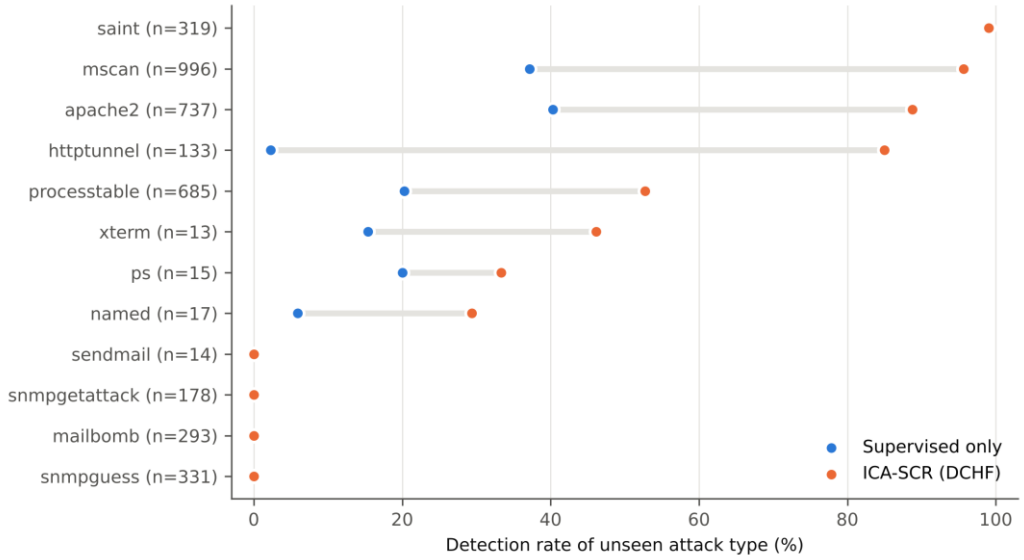


Fig. 6. Detection rate of each unseen NSL-KDD attack type with at least 13 test records, for the supervised channel (blue) and ICA-SCR (orange). Line length shows the gain contributed by the anomaly channel.

5.4 Multi-Class Attribution (E5)

An alert with an attack family will help accelerate triage and playbook selection. The results for ten classes on UNSW-NB15 using the LightGBM channel are presented in Table 8 and Fig. 7. The overall accuracy is 75.7% and the weighted F1 is 0.780. Generic (F1 = 0.985), Reconnaissance (0.866), Normal (0.844) and Exploits (0.705) are recognised reliably. This is because the flow signature of Analysis, Backdoor and DoS overlaps their records, and macro F1 is modest at 0.488, with Analysis and DoS often confused with Exploits (47-52% of their records; Fig. 7). Ablation with balanced class weights was also attempted, which lowered the accuracy to 52.3% due to the shift of benign records into minority classes and hence unweighted training was maintained. These results indicate that, for hunting, it is better to interpret multi-class output as a ranked hypothesis instead of a verdict.

TABLE 8 Per-class multi-class attribution performance on the UNSW-NB15 test partition

Class	**Precision**	**Recall**	**F1-score**	**Support**
Analysis	0.0569	0.1492	0.0824	677
Backdoor	0.0192	0.0875	0.0315	583
DoS	0.5066	0.1220	0.1967	4089
Exploits	0.6385	0.7880	0.7054	11132
Fuzzers	0.2971	0.5793	0.3928	6062
Generic	0.9975	0.9729	0.9850	18871
Normal	0.9591	0.7543	0.8444	37000
Reconnaissance	0.9353	0.8066	0.8662	3496
Shellcode	0.4100	0.7593	0.5325	378
Worms	0.1620	0.5227	0.2473	44
macro avg	0.4982	0.5542	0.4884	82332
weighted avg	0.8353	0.7571	0.7796	82332

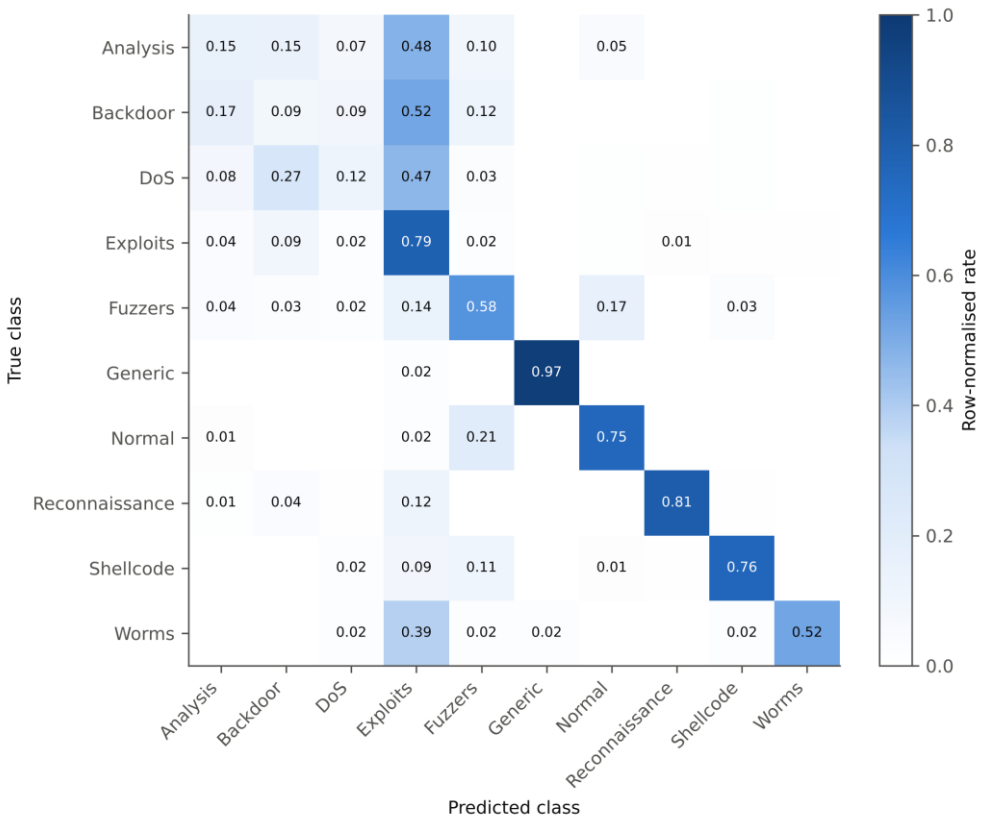


Fig. 7. Row-normalised confusion matrix of the ten-class UNSW-NB15 model. Generic, Reconnaissance, Shellcode and Normal traffic are well separated, whereas Analysis, Backdoor and DoS are often absorbed into Exploits.

5.5 Leave-One-Category-Out Zero-Day Simulation (E4)

TABLE 9. Leave-one-category-out zero-day simulation on UNSW-NB15: detection rate (DR) of the held-out category and FPR on benign test traffic

Held-out category	Test records	Supervised DR (%)	Anomaly-only DR (%)	ICA-SCR DR (%)	Supervised FPR (%)	ICA-SCR FPR (%)
Analysis	677	96.8	84.9	96.8	24.6	27.2
Backdoor	583	100.0	89.5	100.0	26.2	28.2
DoS	4089	99.9	63.3	99.9	26.1	28.1
Exploits	11132	98.3	28.9	98.3	24.9	27.0
Fuzzers	6062	39.0	22.0	40.6	5.6	8.1
Generic	18871	99.8	96.7	100.0	26.2	28.2
Reconnaissance	3496	99.3	9.0	99.3	26.4	28.4
Shellcode	378	98.9	0.8	98.9	26.2	28.2
Worms	44	100.0	15.9	100.0	26.2	28.2

Table 9 and Fig. 8 list a more stringent test in which an entire attack family is removed from training. The supervised channel is still able to detect the majority of the held-out families at very high levels (96.8–100%) on

UNSW-NB15 due to the fact that they share discriminative flow properties, such as anomalous TTL patterns, with the remaining families (Section 5.6). The exception is Fuzzers, which is only found at 39.0% withheld; there ICA-SCR contributes 1.6 points (40.6%). The range of the anomaly channel is 0.8% (Shellcode) to 96.7% (Generic) which validates that low-volume payload-oriented attacks like Shellcode and Worms are virtually undetectable using flow-level anomaly detection. In all heldout families, fusion boosted benign FPR by approximately 2 points (e.g. 26.2% to 28.2%). The Fuzzers run is also enlightening in another respect: the benign FPR for the supervised channel decreases from ~26% to 5.6% when Fuzzers records are withheld. Hence, most of the high false-alarm rate on UNSW-NB15 is due to the benign traffic that mimics the fuzzing traffic, rather than the other families of attacks. In this anomaly set, the anomaly channel is an insurance premium that is seldom paid. In NSL-KDD (Section 5.3), the same mechanism, on the other hand, doubled the detection of unseen types.

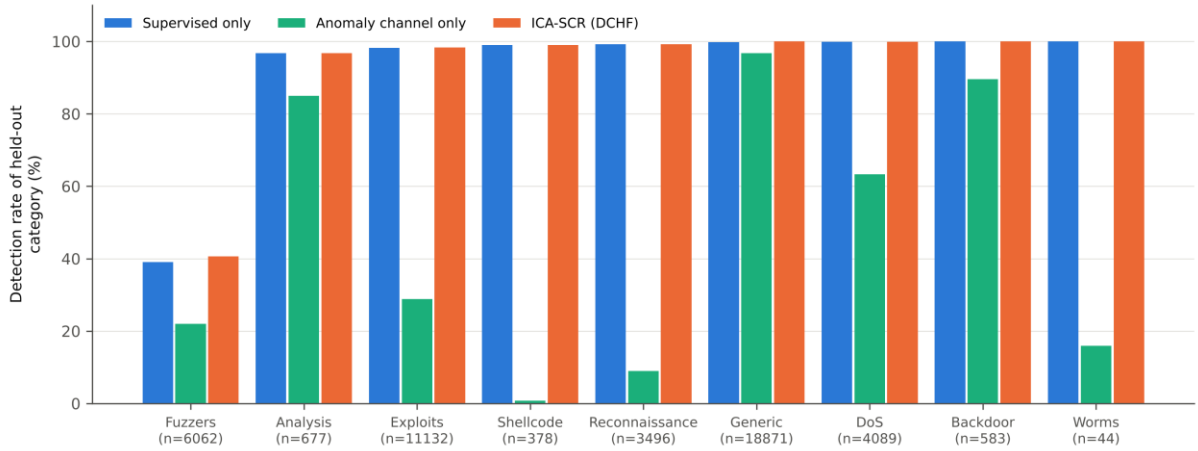


Fig. 8. Leave-one-category-out zero-day simulation on UNSW-NB15. For each held-out attack family, bars show the detection rate of the supervised channel (retrained without that family), the anomaly channel alone and ICA-SCR. Families are ordered by supervised detection rate.

5.6 Explainability (E5)

The top 10 features based on the mean absolute SHAP value of the supervised channel on UNSW-NB15 are presented in Fig. 9 and Table 10. The feature with the largest mean |φ|, which dominates, is source time-to-live (sttl) at 4.18 log-odds units, approximately 5 times the next feature. This is followed by ct_state_ttl (0.86), ct_dst_sport_ltm (0.83), the DNS service indicator (0.74) and source bytes (0.55). These attributions translate to tangible hypotheses for the hunters: if the TTL is unusual, the traffic may be coming from a network distance or tooling that is not expected; if a large number of connections are seen to the same destination and source-port combination, the traffic may be automated scanning or beaconing from a partner segment. The sttl dominance also brings a risk. This is partially a result of the testbed topology in UNSW-NB15 and a model relying on TTL might not be applicable in a partner network. Explainability is thus a validation tool and an analyst aid.

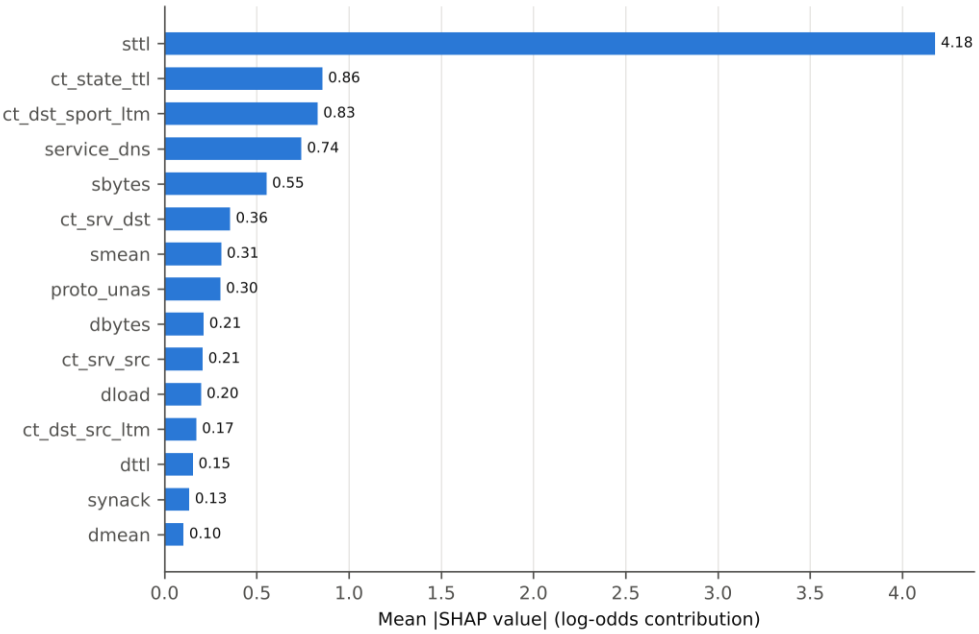


Fig. 9. Global feature importance of the supervised channel on UNSW-NB15 (mean absolute SHAP value over 3,000 test records).

TABLE 10 Top SHAP features and the hunting hypotheses they support

Feature	**Mean |SHAP|**	**Hunting interpretation**
sttl	4.179	Unexpected hop distance or tool-specific TTL; spoofing or foreign infrastructure
ct_state_ttl	0.857	Connection-state/TTL combinations atypical of business flows
ct_dst_sport_ltm	0.830	Repeated destination/source-port pairs: scanning, beaconing
service_dns	0.743	DNS-based reconnaissance, tunnelling or exfiltration
sbytes	0.553	Abnormal upload volume: payload delivery or exfiltration
ct_srv_dst	0.356	Fan-in to one service: brute force or exploitation attempts
smean	0.307	Atypical mean packet size from the source
proto_unas	0.303	Unassigned protocol usage: covert or malformed traffic
dbytes	0.212	Abnormal response volume: data staging or DoS
ct_srv_src	0.207	Fan-out from one source: lateral movement or sweeps

5.7 Hunting Efficiency (E6)

Fig. 10(b) dissects the alerts on NSL-KDD at the operating point. The supervised channel raises 8,323 alerts (8,050 threats and 273 false alarms). ICA-SCR raises 9,774 (9,469 true threats, 305 false alarms). Of the 1451 more alerts, 1419 (97.8%) are true alerts, meaning there's virtually no analyst dilution in coverage with fusion. The results are shown in Fig. 10(a) where all test records are sorted by the priority π assigned by the hunt and the percentage of unseen-type attacks captured is plotted against the increasing alert budget. ICA-SCR is seen to capture slightly more unseen threats at intermediate budgets (66.9% versus 61.2% at k = 10,000). The two curves are otherwise similar, as records identified only by the anomaly channel are ranked after high confidence supervised detections. It is a natural refinement to have a budget-aware ranking that interleaves both channels.

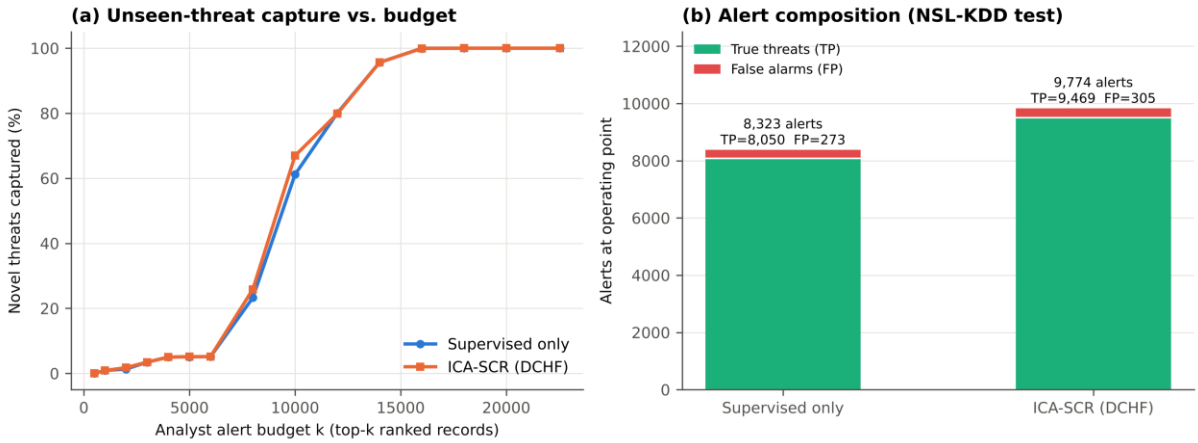


Fig. 10. Hunting efficiency on NSL-KDD: (a) share of unseen-type attacks captured when analysts review the top-k records ranked by priority; (b) alert composition at the operating point. Almost all of the 1,451 additional ICA-SCR alerts are true threats.

5.8 Calibration of the Anomaly Channel

The calibrated anomaly score a_(_x) distribution of the benign and attack test traffic is shown in Fig. 11. On NSL-KDD, attack scores tend to cluster around a high score (above the threshold $\tau = 0.974$), while the benign scores are distributed on [0, 1] as one would expect for a percentile-based score. It can be seen that the benign test traffic is shifted to higher percentiles than the benign training reference on UNSW-NB15, and that a mass of attack traffic is located between the 0.85 percentile and 0.95 percentile, just below $\tau = 0.952$. Both effects account for the lower impact that the anomaly channel has on this dataset and suggest that benign reference calibration (layer L6) is crucial when traffic drifts.

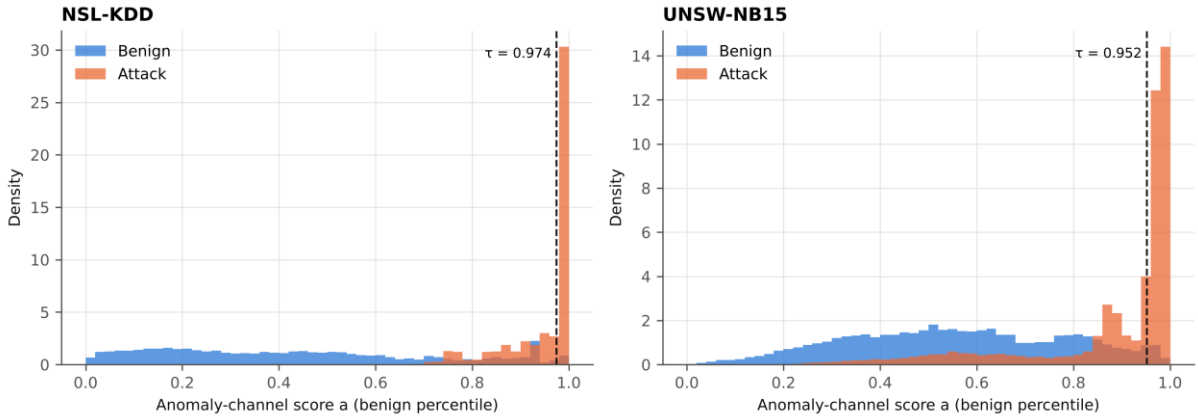


Fig. 11. Distribution of the calibrated anomaly score for benign and attack test records. The dashed line marks the fusion threshold τ (99th benign percentile).

5.9 Scalability and Latency (E6)

TABLE 11 Streaming scalability of the full ICA-SCR scoring pipeline (UNSW-NB15 features, 2 CPU cores)

Micro-batch size	Records scored	Throughput (records/s)	Mean batch latency (ms)	P95 batch latency (ms)	Latency per record (µs)
100	200,000	5,149	19.4	26.0	194.2

1,000	200,000	20,305	49.2	63.2	49.2
10,000	200,000	41,652	240.1	288.6	24.0
50,000	200,000	43,402	1152.0	1239.7	23.0

The end-to-end scoring performance, including both channels, calibration and fusion, is reported in Table 11 for a total of 200,000 streamed records and is plotted in Fig. 12. The throughput increases from approximately 5,100 records/s at 100 records per micro-batch (with a mean latency of 19 ms) to approximately 41,700 records/s at 10,000 records (240 ms) and peaks at 43,400 records/s at 50,000 records. Thus, the pipeline can process approximately 3.7 billion flow records per day on two commodity cores or provide sub-50 ms reaction times for a batch of 1,000 records at approximately 20,000 records/s. This spectrum includes the border traffic for most medium sized supply chain partners and horizontal scaling across partitions is easy as records are scored independently.

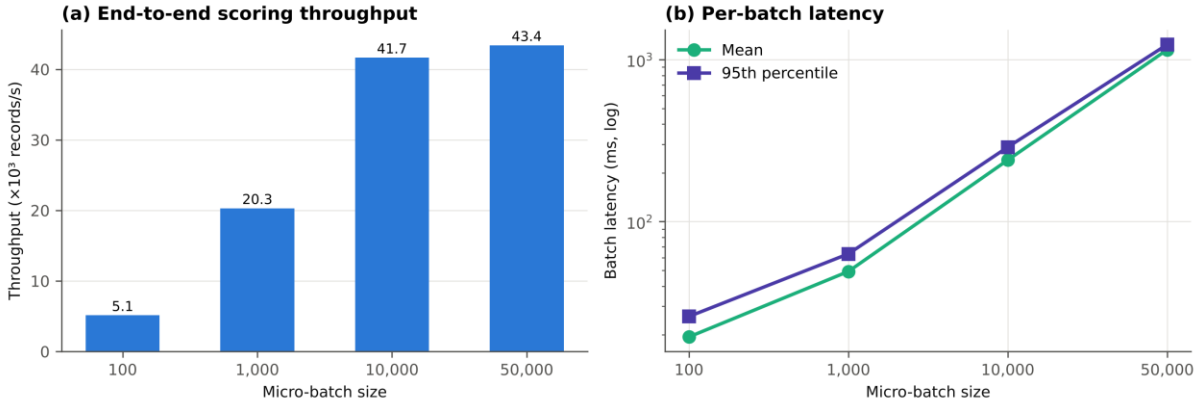


Fig. 12. Scalability of the ICA-SCR scoring pipeline: (a) throughput and (b) mean and 95th-percentile batch latency as a function of micro-batch size.

6. Discussion

6.1 When Does Dual-Channel Analytics Strengthen Resilience?

The two benchmarks provide complimentary, not conflicting, evidence. A benign-calibrated anomaly channel can recover a significant portion of the threats missed by the supervised model when deployment traffic exhibits attack behaviour not seen during training, as with 17 unseen types in NSL-KDD (unseen-type detection increased from 30.3% to 64.5%) with a cost of approximately 0.3 points of FPR. If the new attack families are highly similar to the known attack families, such as UNSW-NB15, a well trained supervised model can already generalise and the anomaly channel will primarily introduce false alarms. The first case is more typical of the supply chain environments. New suppliers and new versions of software, including partner integration, are constantly adding behaviour that is not yet known in the past, while adversaries are intentionally mimicking or evading known behaviour [24]. The results thus indicate that it is appropriate to consider the anomaly channel as calibrated insurance, and that the insurance premium, in the form of a benign quantile q, can be defined directly based on the risk appetence of each segment of partners.

6.2 Operational and Managerial Implications

There are some implications for supply chain security operations. First, there is a need for a shift in evaluation practice. The difference between the cross-validated and shifted-test results (Section 5.2) of 8–20 F1 points illustrates that relying on in-distribution accuracy for procurement or certification decisions could be misleading, and a held-out time period and/or attack family validation should be the standard. Second, the benign-percentile calibration offers a governance lever: security leaders can specify an alert budget, e.g. an additional 1% of benign flows, and audit adherence to this budget, which complements the multi-stakeholder risk management called for in [20] and the collaborative intelligence models of [2] and [3]. Third, SHAP explanations not only reduce triage time, but also expose

fragilities (such as sttl) before they break in production. Lastly, the measured throughput indicates that intelligent analytics is not necessarily big infrastructure, which is important for the smaller suppliers that are often the weakest link in a chain [26] [27]. This aligns with industry data indicating that high adoption of security AI and automation results in shorter breach lifecycles [29].

6.3 Limitations and Threats to Validity

These findings must be interpreted within the context of a number of limitations. (i) Dataset realism. NSL-KDD and UNSW-NB15 are network-level rather than supply chain traces and the mapping to supply chain scenarios in Table 2 is analogical. NSL-KDD is also outdated. Operational partner telemetry, ERP and API logs, and software build events are of high priority for validation. (ii) Feature scope. Content-level attacks (such as snmpguess and sendmail) and software supply chain tampering within build artefacts are not detectable through flow features, and the addition of such payload, log-sequence and SBOM features would increase coverage. (iii) Static calibration, where τ was calculated only from training era benign traffic. Drift takes away the meaning, as seen in fig. 11, online recalibration should be considered. (iv) Simple fusion. The OR rule and the priority function are designed to be transparent. Learned or cost-aware fusion or budget-aware ranking might work better. (v) Single runs. The test-partition results are based on a single test, but the small cross-validation variance (standard deviation ≤ 0.0024) and the significance testing suggest that the results are robust. (vi) Dataset artefacts. Supervised performance in UNSW-NB15 might be overstated due to known artefacts including TTL.

7. Future Work

Future research will (i) assess ICA-SCR for multi-tier supply chain telemetry that accounts for network, ERP or API and CI/CD logs; (ii) incorporate sequence models, such as temporal models from [4] and graph-based learning over supplier interaction graphs to capture multi-stage campaigns; (iii) explore online recalibration and federated training that enable partners to share models without sharing raw data; and (iv) integrate CTI feeds and SBOM intelligence in hypothesis generation for autonomous hunting [5], [8].

8. Conclusion

This paper proposed ICA-SCR, an intelligent cybersecurity analytics framework for supply chain resilience, which combines big-data feature engineering, a supervised gradient-boosting channel, a benign-calibrated autoencoder channel and an isolation-forest channel, Dual-Channel Hunting Fusion, and SHAP-based explanations. In six protocols on two public benchmarks, the framework obtained the highest F1 (0.838) and MCC (0.706) on NSL-KDD, doubled its detection of unseen attack types (30.3% to 64.5%) and obtained 97.8% of its additional alerts to be genuine threats. It also recorded approximately 43,000 records per second when running on two CPU cores. The UNSW-NB15 experiments demonstrated that fusion does not help much when the new attack families are already separable by supervised models, the anomaly channel contributes mostly false alarms. They also demonstrated that the accuracy of a cross-validated model can be up to 8–20 F1 points higher than its actual field performance. All together, the outcomes show a practical design principle for supply chain security operations: combine supervised analytics for known threats with explicitly calibrated anomaly analytics for unknown threats, and balance the two via an auditable number of alerts and clear explanations.

REFERENCES

1. A. A. Osunkanmibi, Y. Adeoye, T. Ogunyankinnu, E. F. Onotole, M. D. Salawudeen, M. A. Abubakar, and A. D. Bello, “Cybersecurity and data protection in supply chains: AI’s role in protecting sensitive financial data across supply chains,” International Journal of Multidisciplinary Research and Growth Evaluation, vol. 6, no. 2, pp. 266–276, 2025.
2. C. Hernández, M. López, J. García, and V. Peterson, “Optimizing collaborative intelligence systems for end-to-end cybersecurity monitoring in global supply chain networks,” International Journal of Responsible Artificial Intelligence, pp. 28–36, 2024.
3. N. Van Hoang, “Human expertise and machine learning in collaborative intelligence frameworks for robust cybersecurity solutions,” Journal of Applied Cybersecurity Analytics, Intelligence, and Decision-Making Systems, vol. 13, no. 12, pp. 1–12, 2023.

4. K. Rathor, B. V. Dhananjayamurthy, A. Jaleel, S. Pal, P. Bhavani, and N. Nishant, "Temporal threat recognition in supply chains: Integrating hidden Markov models for proactive security with AI-driven automated threat hunting," in Proc. 2024 Int. Conf. Advances in Modern Age Technologies for Health and Engineering Science (AMATHE), May 2024, pp. 1–6.
5. N. Sun, M. Ding, J. Jiang, W. Xu, X. Mo, Y. Tai, and J. Zhang, "Cyber threat intelligence mining for proactive cybersecurity defense: A survey and new perspectives," IEEE Communications Surveys & Tutorials, vol. 25, no. 3, pp. 1748–1774, 2023.
6. S. Anne, "Artificial intelligence and cybersecurity integration: Building resilient digital ecosystems in the digital age: A comprehensive review," Journal of Computational Systems & Engineering Insights, vol. 3, no. 1, pp. 11–24, 2025.
7. N. A. F. Shakil, R. Mia, and I. Ahmed, "Applications of AI in cyber threat hunting for advanced persistent threats (APTs): Structured, unstructured, and situational approaches," Journal of Applied Big Data Analytics, Decision-Making, and Predictive Modelling Systems, vol. 7, no. 12, pp. 19–36, 2023.
8. S. R. Sindiramutty, "Autonomous threat hunting: A future paradigm for AI-driven threat intelligence," arXiv preprint arXiv:2401.00286, 2024.
9. I. H. Sarker, "Introduction to AI-driven cybersecurity and threat intelligence," in AI-Driven Cybersecurity and Threat Intelligence: Cyber Automation, Intelligent Decision-Making and Explainability. Cham, Switzerland: Springer Nature Switzerland, 2024, pp. 3–19.
10. A. Raji, A. Olawore, A. Mustapha, and J. Joseph, "Integrating artificial intelligence, machine learning, and data analytics in cybersecurity: A holistic approach to advanced threat detection and response," World Journal of Advanced Research and Reviews, vol. 20, no. 3, pp. 2005–2024, 2023.
11. V. Nagamalla and R. K. Sanapala, "Integrating predictive big data analytics with behavioral machine learning models for proactive threat intelligence in industrial IoT cybersecurity," International Journal of Wireless & Ad Hoc Communication, vol. 7, no. 2, p. 8, 2023.
12. P. S. S. R. Patchamatla, "Developing multi-agent AI-enabled cloud-native cyber defense framework for intelligent threat hunting and response," International Journal of Research and Applied Innovations, vol. 8, no. 6, pp. 13218–13229, 2025.
13. A. Manoharan and M. Sarker, "Revolutionizing cybersecurity: Unleashing the power of artificial intelligence and machine learning for next-generation threat detection," 2023, doi: 10.56726/IRJMETS32644.
14. C. C. Okpala and E. O. Chukwumuanya, "The future of cybersecurity: Predictive analytics and machine learning applications," Journal of Engineering Research and Applied Science, vol. 14, no. 2, pp. 190–201, 2025.
15. S. Ghasemshirazi and G. Shirvani, "Securing the future: Proactive threat hunting for sustainable IoT ecosystems," arXiv preprint arXiv:2406.14804, 2024.
16. H. Volkova, "Platformization in the era of AI-driven cyber warfare: Integrating deep learning, transfer learning, and behavioral analytics for enterprise resilience," Library of Ethiopian International Journal of Multidisciplinary Research, vol. 12, no. 10, pp. 854–859, 2025.
17. E. Zhang, M. Chen, S. Wang, and D. Liu, "Artificial intelligence for cybersecurity analytics: A review of deep learning, hybrid architectures, and operational resilience," Journal of AI Analytics and Applications, vol. 3, no. 2, pp. 1–27, 2025.
18. D. Chinnathambi, D. Dinesh, M. Shanmugavel, and V. Shanmugadass, "Securing sustainable supply chains through blockchain and AI integration," in Sustainable Supply Chains and Carbon Footprint Reduction: The Blockchain Advantage, 2025, pp. 221–235.
19. W. Richard and W. K. Akram, "Big data's role in supply chain optimization: Enhancing e-commerce with scalable data storage and future-proof cyber defense strategies," ResearchGate Journal of Business & Technology, vol. 30, no. 2, pp. 98–115, 2024.
20. L. Rojas, A. G. Sánchez, P. A. Cruz, C. B. Mendoza, C. V. Peterson, D. F. Gómez, and J. R. Vargas, "Evaluation of collective intelligence mechanisms for multi-stakeholder cybersecurity risk management in supply ecosystems," QuestSquare, vol. 7, pp. 55–66, 2024.
21. S. S. Abosuliman, "Deep learning techniques for securing cyber-physical systems in supply chain 4.0," Computers & Electrical Engineering, vol. 107, Art. no. 108637, 2023.
22. M. Shakir, "Smart cyber intelligence and machine learning models for secure cloud-native enterprise data platforms," International Journal of Science, Research and Technology, vol. 7, no. 6, pp. 13273–13283, 2024.
23. A. Aggarwal, "Adoption of AI-powered threat intelligence in cloud infrastructures for real-time risk mitigation and automated anomaly detection through predictive security analytics," Journal of Digital Security and Forensics, vol. 2, no. 1, pp. 139–150, 2025.
24. Z. Tan, S. P. Parambath, C. Anagnostopoulos, J. Singer, and A. K. Marnerides, "Advanced persistent threats based on supply chain vulnerabilities: Challenges, solutions, and future directions," IEEE Internet of Things Journal, vol. 12, no. 6, pp. 6371–6395, 2025.
25. S. Vengathattil and S. M. Shaffi, "Advanced network security through predictive intelligence: Machine learning approaches for proactive threat detection—an experimental study," Methodology, vol. 2, no. 3, 2025.
26. S. R. Sindiramutty, C. E. Tan, W. W. Goh, S. Balakrishnan, N. Hamzah, and R. Akbar, "Securing the supply chain: Cybersecurity strategies for logistics resilience," in Navigating Cyber Threats and Cybersecurity in the Logistics Industry. IGI Global Scientific Publishing, 2024, pp. 300–365.
27. S. Herath, J. G. Johannsson, Y. R. Shrestha, and G. von Krogh, "Supply chain resilience through artificial intelligence-enabled predictive maintenance: A design science research study," SSRN, Paper no. 5153658, 2025.

28. Verizon Business, “2025 Data Breach Investigations Report,” Verizon, Basking Ridge, NJ, USA, Apr. 2025. [Online]. Available: https://www.verizon.com/business/resources/reports/dbir/
29. IBM Security and Ponemon Institute, “Cost of a Data Breach Report 2025,” IBM Corp., Armonk, NY, USA, Jul. 2025. [Online]. Available: https://www.ibm.com/reports/data-breach
30. S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” in Advances in Neural Information Processing Systems, vol. 30, 2017, pp. 4765–4774.
31. G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, “LightGBM: A highly efficient gradient boosting decision tree,” in Advances in Neural Information Processing Systems, vol. 30, 2017, pp. 3146–3154.
32. F. T. Liu, K. M. Ting, and Z.-H. Zhou, “Isolation forest,” in Proc. 8th IEEE Int. Conf. Data Mining (ICDM), Pisa, Italy, 2008, pp. 413–422.
33. M. Tavallaee, E. Bagheri, W. Lu, and A. A. Ghorbani, “A detailed analysis of the KDD CUP 99 data set,” in Proc. IEEE Symp. Computational Intelligence for Security and Defense Applications (CISDA), Ottawa, ON, Canada, 2009, pp. 1–6.
34. N. Moustafa and J. Slay, “UNSW-NB15: A comprehensive data set for network intrusion detection systems (UNSW-NB15 network data set),” in Proc. Military Communications and Information Systems Conf. (MilCIS), Canberra, ACT, Australia, 2015, pp. 1–6.