

Artificial Intelligence for secure cloud computing in Intrusion detection systems and emerging research trends: A Systematic Review

Mohd Asgar Ali

School of Computer Application, Lovely Professional University, Phagwara, Punjab, India

Aliasgarali05@gmail.com

Amanpreet Singh

School of Computer Application, Lovely Professional University, Phagwara, Punjab, India

Apsj24@gmail.com

Abstract

Background: Cloud computing's elasticity, multi-tenancy, and distributed service model (IaaS/PaaS/SaaS) create an attack surface that traditional signature-based intrusion detection systems (IDS) cannot adequately cover. Artificial intelligence (AI) — spanning classical machine learning (ML), deep learning (DL), federated learning, graph neural networks (GNNs), and transformer architectures — has become the dominant approach for cloud-native intrusion detection, but the resulting literature is fragmented across technique families, datasets, and evaluation protocols.

Objective: This paper systematically reviews AI-driven IDS research for cloud computing published between 2021 and 2026, guided by five research questions covering technique evolution, dataset adequacy, performance trade-offs, explainability, and robustness/adaptability gaps.

Methods: We followed the PRISMA 2020 methodology across six academic databases, applying explicit inclusion/exclusion criteria and a five-dimension quality-assessment rubric. From an initial pool of candidate records drawn from the authors' collected bibliography and targeted database searches, we retained a verified core set of quantitatively-reported cloud-IDS studies for detailed comparative synthesis, supplemented by a broader set of topically-relevant studies (2021–2026) used to characterize trend and taxonomy coverage.

Results: We propose a multidimensional AI-IDS taxonomy (detection technique × AI paradigm × deployment scope × cloud service model), a dataset benchmark matrix spanning legacy (KDD99, NSL-KDD) through cloud-and-IoT-specific benchmarks (UNSW-NB15, CSE-CIC-IDS2018, Bot-IoT, CICIoT2023, TON_IoT, 5G-NIDD), and a technique-by-technique comparative synthesis covering classical ML/ensemble, CNN-based, RNN/LSTM/GRU-based, hybrid deep learning, federated learning, and metaheuristic-optimized IDS models. Reported accuracies for cloud-targeted models range from roughly 92% to above 99.9%, but comparability across studies is undermined by inconsistent datasets, non-standardized metrics, and near-absent explainability reporting.

Conclusions: Federated learning and metaheuristic feature selection are maturing rapidly for cloud IDS, but explainable AI (XAI), graph-based modeling, transformer/LLM-assisted detection, concept-drift handling, and adversarial robustness remain substantially under-explored in the cloud-specific literature

relative to the broader network-IDS and cybersecurity-XAI literature. We close with a five-pillar research roadmap — XAI-integrated cloud IDS, federated IDS, GNN-based IDS, transformer/LLM-assisted SOC and cloud IDS, and adversarial-robustness/concept-drift-aware IDS — intended to guide the next generation of cloud-native intrusion detection research.

Keywords

Cloud Computing; Intrusion Detection System; Artificial Intelligence; Machine Learning; Deep Learning; Systematic Review; PRISMA; Explainable AI; Federated Learning; Graph Neural Networks; Adversarial Robustness.

1. Introduction

Cloud computing has become the default infrastructure model for modern IT, offering on-demand, elastic, and cost-efficient resources across the Infrastructure-, Platform-, and Software-as-a-Service (IaaS/PaaS/SaaS) layers [1]. This shift concentrates enormous volumes of sensitive data and critical workloads behind shared, virtualized, multi-tenant infrastructure, which in turn expands the attack surface available to adversaries: insecure APIs, hypervisor-level exploits, lateral movement between co-located tenants, and large-scale distributed denial-of-service (DDoS) campaigns are all cloud-amplified variants of long-standing network threats [2] [11].

Intrusion Detection Systems (IDS) remain the primary automated defense for identifying malicious activity within this environment, whether deployed as network-based IDS (NIDS) at cloud ingress/egress points, host-based IDS (HIDS) on individual virtual machines, or distributed IDS (DIDS) spanning multiple availability zones [5] [6]. However, classical signature-based detection is structurally unsuited to the cloud: it cannot scale elastically with tenant demand, cannot easily distinguish between legitimate multi-tenant traffic and cross-tenant anomalies, and is blind to zero-day and low-and-slow attacks such as Advanced Persistent Threats (APTs) [3] [9]. These limitations have driven a decade-long shift toward AI-based detection, in which machine learning and deep learning models learn traffic and behavioral patterns directly from data rather than from hand-authored signatures [6] [19].

This shift has produced a large and fast-growing body of literature, but one that is fragmented in three specific ways relevant to practitioners and researchers alike. First, techniques are rarely compared across families — CNN-based, LSTM-based, ensemble, federated, and metaheuristic-optimized IDS models are usually evaluated only against earlier instances of the same family, not against the broader technique landscape. Second, dataset usage is inconsistent: studies mix decades-old benchmarks (KDD99) with recent, cloud-native traffic captures (CICIoT2023, TON_IoT, 5G-NIDD) without a shared understanding of which benchmarks actually represent cloud conditions such as multi-tenancy and elasticity. Third, and most acutely, explainability and adversarial robustness — both extensively studied in the general cybersecurity-XAI literature [16] [17] — are almost entirely absent from the cloud-IDS literature specifically, despite being flagged repeatedly as open problems.

1.1 Related work and contributions

Several prior surveys have examined AI-based intrusion detection, but none combine a transparent systematic (PRISMA-based) selection process with a cloud-specific scope and coverage of the newest technique families (federated learning, GNNs, transformers, XAI). Table 1 positions this survey against three representative prior works drawn from the same literature base.

Table 1. How this survey advances prior cloud-/AI-IDS surveys.

Survey (Year)	Scope	Systematic (PRISMA) method	XAI covered	Federated / GNN / Transformer	Adversarial robustness	Research questions stated
Nassif et al. (2021) [19], "Machine Learning for Cloud Security: A Systematic Review"	ML for cloud security, broad (not IDS-specific)	Partial	–	–	–	–
Rana et al. (2022) [6], "Intrusion Detection Systems in Cloud Computing Paradigm"	Cloud IDS, analysis/overview	–	–	–	–	–
Chinnasamy et al. (2025) [14], DL-driven NIDS systematic review	DL for network IDS (not cloud-exclusive)	Partial	–	Partial (transformer-adjacent)	–	–
CloudGuard AI (this work, 2026)	AI-IDS specifically in cloud computing	Full PRISMA 2020	✓	✓	✓	✓ (5 RQs)

Note: rows for the three prior surveys reflect their stated scope as indexed; a full methodological audit of each (search strategy, inclusion criteria, study count) is recommended before this table is presented as final in a submission-ready manuscript, and is flagged here as an open verification item rather than asserted as complete.

Building on this gap analysis, the contributions of this work are:

- A transparent, PRISMA 2020-based selection methodology for cloud-specific AI-IDS literature (2021–2026), including explicit inclusion/exclusion criteria and a five-dimension quality-assessment rubric.
- A multidimensional AI-IDS taxonomy that cross-references detection technique, AI paradigm, deployment strategy, and cloud service model (IaaS/PaaS/SaaS), extending prior single-axis taxonomies.
- A dataset benchmark matrix spanning legacy network-IDS datasets through 2023–2025 cloud- and IoT-native benchmarks, annotated for cloud-representativeness and known limitations.
- A technique-by-technique comparative synthesis (classical ML/ensemble, CNN, RNN/LSTM/GRU, hybrid DL, federated learning, metaheuristic-optimized models) following a define-mechanism → review-studies → compare → strengths → weaknesses → gap structure.
- A research-gap matrix and five-pillar future roadmap covering XAI, federated IDS, GNN-based IDS, transformer/LLM-assisted detection and SOC integration, and adversarial-robustness/concept-drift handling — areas that are mature in the general AI-security literature but largely unaddressed for cloud-specific IDS.

1.2 Research Questions

- RQ1 — Which AI/ML/DL paradigms (classical ML, deep learning, ensemble, federated, graph-based, transformer-based) have been applied to intrusion detection in cloud environments between 2021 and 2026, and how has technique adoption shifted over this period?
- RQ2 — Which datasets are used to train and evaluate cloud-targeted AI-IDS, and to what extent do they represent real multi-tenant, elastic, distributed cloud traffic as opposed to generic or IoT-only traffic?
- RQ3 — What detection-performance trade-offs (accuracy, false-alarm rate, training/inference latency, computational cost) characterize each AI-IDS technique family when applied at cloud scale?
- RQ4 — How explainable and trustworthy are current cloud AI-IDS models, and to what extent has explainable AI (XAI) been integrated into cloud-specific intrusion detection, as distinct from generic network IDS?
- RQ5 — What robustness, privacy, and adaptability gaps — adversarial evasion/poisoning, concept drift, cross-tenant data leakage, regulatory constraints — currently limit real-world deployment of AI-IDS in production cloud environments?

1.3 Paper selection methodology (PRISMA 2020)

This review follows the PRISMA 2020 reporting framework, structured into identification, screening, eligibility, and inclusion stages.

1.3.1 Identification

Candidate records were drawn from IEEE Xplore, ACM Digital Library, ScienceDirect (Elsevier), SpringerLink, Google Scholar, and Scopus, using combinations of the search terms "cloud computing" OR "cloud security" AND "intrusion detection" OR IDS AND "machine learning" OR "deep learning" OR "artificial intelligence" OR "federated learning" OR "explainable AI", restricted to publications dated 2021–2026. This search was supplemented with the authors' own working bibliography of cloud/IDS/AI-security literature compiled during the initial drafting of this review.

1.3.2 Screening and eligibility

Title/abstract screening removed duplicates and clearly out-of-scope records (e.g., records addressing cloud security controls with no AI-IDS component, or AI/ML papers with no security application). Full-text screening was then applied against the inclusion/exclusion criteria below.

Inclusion criteria:

- Peer-reviewed journal or conference publication (or a clearly high-impact preprint, per the convention used in comparable XAI-cybersecurity surveys), dated 2021–2026.
- Explicitly addresses intrusion, anomaly, or attack detection in a cloud, fog, or cloud-adjacent (IoT-to-cloud, mobile-cloud, MEC) environment.
- Uses an AI/ML/DL technique as the core detection mechanism (not purely rule-based or purely cryptographic defenses).

- Reports at least one quantitative evaluation metric (accuracy, precision, recall, F1, false-alarm rate, AUC, or equivalent).

Exclusion criteria:

- Generic network-IDS papers with no cloud-specific framing or evaluation context.
- Purely conceptual, position, or vision papers with no empirical evaluation.
- Duplicate or extended-abstract versions of an already-included full paper.
- Papers whose subject matter is unrelated to cloud/AI/IDS despite superficial keyword overlap (one such record — an enzymatic/biochemical kinetics study erroneously present in an earlier draft of this bibliography — was identified during this screening step and removed).

1.3.3 Quality assessment

Each included study was scored 0–2 on five dimensions, for a maximum of 10 points, adapting the mixed quantitative/qualitative rubric used in prior XAI-cybersecurity survey work [13]: (a) dataset appropriateness for the claimed threat model; (b) evaluation rigor (held-out test set or cross-validation, as opposed to train-set-only reporting); (c) presence of baseline/comparative evaluation; (d) reproducibility (availability of code, data, or sufficient implementation detail); (e) explicit cloud-specificity of the deployment scenario. Studies scoring below 4/10 were flagged for exclusion from the comparative synthesis tables in Section 6, though they may still be referenced narratively where topically relevant.

1.4 Paper organization

Section 2 presents a multidimensional taxonomy of AI-IDS in cloud computing. Section 3 reviews available datasets as a benchmark matrix. Section 4 provides AI/ML/DL background. Section 5 reviews recent studies technique-by-technique. Section 6 presents the comparative study table. Section 7 discusses attacks and threats against cloud AI-IDS itself. Section 8 presents the challenges and research-gap matrix. Section 9 proposes a five-pillar future research roadmap. Section 10 concludes.

2. AI-IDS Taxonomy in Cloud Computing

IDS for cloud computing can be categorized along several largely-independent axes: detection technique, deployment strategy, scope of protection, response strategy, cloud service model, and — the axis most relevant to this review and typically absent from prior cloud-IDS taxonomies — the underlying AI paradigm. Table 2 summarizes these axes; a given system in the literature is typically describable as a point in this multidimensional space (e.g., "an anomaly-based, network-deployed, IaaS-targeted IDS using a federated CNN").

Table 2. Multidimensional AI-IDS taxonomy for cloud computing.

Axis	Categories	Notes
Detection technique	Signature-based; Anomaly-based; Hybrid	Signature-based struggles with zero-day/APT; anomaly-based trades false positives for novel-attack coverage [11]

Axis	Categories	Notes
Deployment strategy	Network-based (NIDS); Host-based (HIDS); Distributed (DIDS)	DIDS increasingly required for multi-availability-zone cloud deployments [12]
Scope of protection	System-level; Application-level; Data-centric	Data-centric IDS is emerging as multi-tenant data-exfiltration risk grows [13]
Response strategy	Active (auto-mitigation); Passive (alert-only)	Active response raises risk of automated false-positive disruption in shared infrastructure
Cloud service model	IaaS; PaaS; SaaS	Each layer has a distinct threat surface: VM/hypervisor (IaaS), app deployment (PaaS), user/application transactions (SaaS) [14] [15]
AI paradigm (this review's added axis)	Classical ML/ensemble; CNN-based; RNN/LSTM/GRU-based; Hybrid DL; Federated learning; Metaheuristic-optimized; Graph-based (GNN); Transformer/attention-based; XAI-integrated	Distribution and maturity of each category is reviewed in Section 5

Dataset	Year	Native environment	Cloud-representative?	Key limitation
KDD99 / KDD CUP 99	1999	General network (DARPA98-derived)	No	Outdated, highly skewed targets, redundant patterns, non-stationary train/test split [17]
NSL-KDD	2009	General network (KDD99-derived)	No	Improved balance over KDD99 but still lacks low-footprint/APT representation [17]
DEFCON-8 / DEFCON-10	2000/2002	CTF competition traffic	No	Overwhelmingly attack traffic; not representative of real background traffic [18]
CAIDA (OC48 / DDoS / 2016 traces)	2002–2016	Internet backbone	Limited	Event-specific, anonymized, lacks ground truth for many records
LBNL	2005	Enterprise network (header-only)	No	No labeling process; header-only, no engineered features
UNSW-NB15	2015	Synthetic (IXIA PerfectStorm)	Partial	Synthetic generation environment; 10-class (1 benign, 9 attack) coverage [23]
CICIDS2017 / CSE-CIC-IDS2018	2017/2018	Simulated enterprise + AWS cloud	Yes (2018 variant)	Profiling-based generation; ground truth not fully released [21]
TUIDS	~2012	University testbed	No	Flow-level data lacks novel features beyond the capture process
Bot-IoT	2019	Cyber Range Lab, UNSW Canberra	Yes (IoT-cloud)	Botnet-centric; complex multi-format source files [24]
IoTPoT	2015	Honeypot capture	Limited	Restricted view (honeypot-only); Telnet-focused [25]
N-BaloT	2018	IP-camera + IoT testbed (9 devices, Mirai-infected)	Yes (IoT-cloud)	Narrow device set; attack diversity limited to 9 categories

Dataset	Year	Native environment	Cloud-representative?	Key limitation
CICIoT2023	2023	IoT-to-cloud testbed	Yes	Recent, large, but adoption/benchmarking still maturing
TON_IoT	2021	IoT/IIoT + cloud telemetry	Yes	Heterogeneous sub-datasets complicate cross-study comparison
5G-NIDD	2022	5G network testbed	Partial (5G-cloud edge)	5G-specific; limited generalization to non-5G cloud contexts
FLNET2023	2023	Federated-learning network simulation	Yes	Purpose-built for FL evaluation; newer, less independently validated

This taxonomy highlights the diversity and complexity of IDS solutions tailored for cloud environments, and — critically for the synthesis in Section 5 — makes explicit that most existing surveys collapse the AI-paradigm axis into a single "ML vs. DL" split, obscuring meaningful differences between, for example, federated and non-federated deep learning approaches to the same detection task.

3. Dataset Benchmark Matrix

Evaluating any IDS requires a reliable dataset containing representative benign and malicious activity. Table 3 summarizes the datasets used across the reviewed cloud-IDS literature, from legacy general-network benchmarks through recent cloud- and IoT-native captures.

Table 3. Dataset benchmark matrix.

4. AI, Machine Learning, and Deep Learning Background

AI is an umbrella term for systems that emulate aspects of human intelligence; within cybersecurity, the dominant sub-fields are classical machine learning (ML), deep learning (DL), and — increasingly — federated and graph-based learning. All learn a mapping from network/system observations to a benign/malicious label (or attack-class label), differing chiefly in how that mapping is represented and learned.

4.1 Machine learning

Classical ML for IDS spans supervised approaches (Logistic Regression, SVM, Decision Trees, Random Forest, Naïve Bayes, KNN), which learn from labeled normal/malicious traffic, and unsupervised approaches (K-Means, SOM), which extract structure from unlabeled data. Ensembles combining multiple base learners are common and often outperform any single model, at the cost of increased training/inference overhead and, in several reviewed studies, sensitivity to overfitting on synthetic datasets [11].

4.2 Deep learning

DL models — CNNs, RNNs/LSTMs/GRUs, autoencoders, and their hybrids — automatically learn hierarchical feature representations from raw or lightly-processed traffic, removing the need for manual

feature engineering. This is particularly valuable in cloud settings where traffic volume and dimensionality make manual feature design impractical. The trade-off is substantial computational cost, large labeled-data requirements, and — most relevant to RQ4 — a "black-box" character that complicates operational trust and incident-response justification [6].

4.3 Emerging paradigms relevant to cloud IDS

Federated learning (FL) allows multiple tenants or edge nodes to collaboratively train a shared detection model without centralizing raw traffic, directly addressing the multi-tenancy privacy constraint identified in Section 1. Graph Neural Networks (GNNs) model network traffic or system-call relationships as graphs, well-suited to botnet and lateral-movement detection. Transformer/attention-based architectures, borrowed from NLP, are beginning to appear in traffic-classification and packet-sequence modeling. Metaheuristic optimization (e.g., squirrel search, artificial bee colony, Tasmanian devil optimization) is frequently paired with DL models in the reviewed literature to perform feature selection or hyperparameter tuning under cloud-scale dimensionality constraints. Each of these paradigms is reviewed with its associated studies in Section 5.

A typical AI-IDS development workflow proceeds through data collection, data processing (cleaning, normalization, feature engineering), model selection, training, evaluation (held-out or k-fold), and hyperparameter tuning, though the specific sequence varies by technique family and task.

5. Recent AI-IDS Studies in Cloud Computing

Following the pattern established in comparable systematic reviews of AI techniques in cybersecurity [13], each subsection below defines the technique category, explains its mechanism, reviews the grouped studies together, compares their reported findings, and then states strengths, weaknesses, and the unresolved gap for that category.

5.1 Classical ML and ensemble-based cloud IDS

Mechanism: classical ML/ensemble approaches apply dimensionality-reduction and feature-selection techniques ahead of a shallow classifier (or a combination of classifiers) to distinguish normal from malicious cloud traffic, typically favored for lower training cost relative to deep learning.

Bai and Sudha proposed combining Principal Component Analysis (PCA) and Singular Value Decomposition (SVD) for dimensionality reduction ahead of a deep-learning classifier evaluated on UNSW-NB15 for cloud-forensics intrusion detection, reporting that the combined pipeline outperformed comparison approaches at identifying fraudulent data streams in cloud-based networks [30] [31]. Ioannou et al. proposed a three-stage Medical-IoT (MIoT) pipeline combining an Enhanced (ensemble) Random Forest first-stage classifier — reported at 99.98% accuracy — with a One-Class SVM second-stage anomaly detector reported at 99.7% accuracy, coordinated by a Federated Learning model-update stage, validated on a Raspberry-Pi MIoT gateway testbed [41]. Attou et al. combined Random Forest-based feature selection with a Radial Basis Function Neural Network (RBFNN) classifier, evaluated on Bot-IoT and NSL-KDD, reporting accuracy

above 92% using a minimal feature subset and an increase in Matthew's Correlation Coefficient from 28% to 93% on the imbalanced Bot-IoT set [44].

Comparative findings: across these three studies, ensemble/classical pipelines reach very high reported accuracy (92–99.98%) but on markedly different datasets and threat models (cloud forensics vs. MIoT vs. Bot-IoT/NSL-KDD), which limits direct cross-study comparability — a recurring pattern across nearly all technique categories reviewed here (see Section 8).

Strengths: low computational cost relative to deep learning; strong performance on well-curated, moderate-dimensionality datasets; RF-based feature selection is a natural pairing for cloud environments with high feature-dimensionality traffic.

Weaknesses: sensitivity to class imbalance without explicit correction (as evidenced by Attou et al.'s low baseline MCC before optimization); limited capacity to model complex sequential/temporal attack patterns relative to DL; reliance on hand-selected or wrapper-based feature sets that may not generalize across cloud providers.

Unresolved gap: no reviewed classical-ML cloud-IDS study reports explainability analysis of the resulting feature-importance rankings, despite feature importance being a natural, low-cost explainability signal for tree-based and RF-based models specifically.

5.2 CNN-based cloud IDS

Mechanism: convolutional architectures, originally developed for image processing, are applied to cloud IDS either directly on tabular/flow features or on traffic reshaped into image-like representations (e.g., system-call sequences rendered as images).

Vibhute and Nakum proposed a deep CNN for near-real-time network intrusion detection in an imbalanced cloud environment, using a Random Forest model to select the most relevant input features, evaluated on CSE-CIC-IDS2018, and reported 97.07% testing accuracy (2.93% error rate), with precision, recall, and F1-score of 98.11%, 96.93%, and 97.52% respectively [27]. Joraviya et al. proposed DL-HIDS, converting system-call logs into image representations for CNN classification in containerized cloud environments, evaluated on the Leipzig Intrusion Detection DataSet-2019, and reported improvements over prior system-call-based methods across accuracy, precision, recall, F1, and false-positive rate, while noting that determining the optimal image feature size remains an open, largely empirical, design choice [28]. Jansi Sophia Mary and Mahalakshmi proposed a Human-In-The-Loop (HITL) cloud IDS using a CNN classifier and a non-zero-sum game-theoretic evaluation model, reporting successful detection across five attack classes with claimed outperformance of comparable literature approaches [29].

Comparative findings: CNN-based cloud IDS in the reviewed set consistently pairs the convolutional classifier with an upstream feature-selection or representation-transformation step (RF-based feature selection, image encoding of system calls), suggesting that raw high-dimensional cloud traffic is rarely fed directly into a CNN without dimensionality reduction.

Strengths: strong pattern-recognition performance on transformed/curated inputs; the image-encoding approach (Joraviya et al.) sidesteps redundant system-call/metadata features that degrade detection rate in prior pretrained-CNN approaches; near-real-time feasibility demonstrated at CSE-CIC-IDS2018 scale.

Weaknesses: image-size and encoding-scheme selection is empirical and not yet systematized; HITL designs (Jansi Sophia Mary and Mahalakshmi) introduce human-response latency that may be incompatible with elastic, high-throughput cloud traffic; none of the three studies reports Grad-CAM or comparable visual explainability despite CNNs being the technique family for which such methods are best developed in the general XAI literature [12].

Unresolved gap: transfer of CNN-specific XAI techniques (Grad-CAM and successors, well-established for malware-image classification in the broader cybersecurity-XAI literature) into cloud-IDS-specific image-encoded representations (e.g., Joraviya et al.'s system-call images) is, to this review's knowledge, unaddressed.

5.3 RNN, LSTM, and GRU-based cloud IDS

Mechanism: recurrent architectures model the temporal/sequential structure of network flows or system events, capturing multi-step attack patterns that flow-level, non-sequential models miss.

Kavitha et al. proposed a filter-wrapper-embedded ensemble feature-selection stage (FEFS) feeding a hybrid RNN model optimized via Tasmanian Devil Optimization (TDO), implemented in MATLAB and compared against a conventional RNN, DNN, and an RNN-genetic-algorithm hybrid across sensitivity, F-measure, precision, recall, and accuracy [43]. Hasimi et al. proposed a feed-forward Artificial Neural Network (ANN) for cloud security, using a Kaggle-sourced dataset, and focused primarily on the practical integration steps (data preparation cost, model development, ongoing monitoring) required to deploy ANN-based detection in an operational cloud-security strategy rather than on raw accuracy maximization [36].

Comparative findings: relative to the CNN and classical-ML categories above, this group of studies is comparatively thin on standardized quantitative benchmarking (Kavitha et al. compares against internally-defined RNN variants rather than external published baselines; Hasimi et al. does not report head-to-head accuracy against alternative technique families), which itself is a finding — recurrent-architecture cloud IDS work in the reviewed sample skews toward integration feasibility and metaheuristic-optimization framing rather than pure benchmark-accuracy competition.

Strengths: metaheuristic-optimized RNN variants (TDO-RNN) directly address the hyperparameter-tuning burden that is a known DL deployment cost in resource-constrained cloud/fog tiers; feed-forward ANN approaches are comparatively lightweight and easier to operationalize than deeper recurrent stacks.

Weaknesses: comparison sets in the reviewed RNN/LSTM/GRU studies are narrower than in the CNN category, limiting confidence in generalizability claims; neither reviewed study reports concept-drift behavior over time, despite recurrent models' natural fit for temporal drift monitoring.

Unresolved gap: no reviewed recurrent-architecture cloud-IDS study evaluates performance degradation under evolving/drifting traffic distributions — a gap directly relevant to RQ5 and elaborated in Section 9.

5.4 Hybrid deep learning and metaheuristic-optimized cloud IDS

Mechanism: hybrid approaches combine two or more DL/optimization components — most commonly a metaheuristic feature-selection or hyperparameter-search stage feeding a deep classifier — to jointly address the dimensionality and tuning burdens that pure single-model DL approaches face at cloud scale.

Yaras and Dener designed a hybrid one-dimensional CNN-LSTM model, trained and tested using PySpark/Apache Spark in a Google Colaboratory environment on the CICIoT2023 and TON_IoT datasets, with features reduced via correlation-based selection; compared against ten ML/DL baselines, the model reached 99.995% accuracy for binary classification and 99.96% for multiclass classification on CICIoT2023, and 98.75% binary-classification accuracy on TON_IoT [46]. Sarkar, Keserwani, and Govil combined an Improved Squirrel Search Algorithm (ISSA) for feature extraction with a Modified Deep Belief Network (MDBN) for anomaly classification on UNSW-NB15, reporting a reduction in training time from 103.51 to 0.108 seconds and testing time from 29.62 to 0.054 seconds relative to baseline, alongside 99.8% accuracy and a false-alarm rate of 0.02% — the strongest reported combination of accuracy and computational efficiency in the reviewed sample [32]. Laassar et al. proposed a Binary Artificial Bee Colony Network (BABCN) combining depth-first-search-improved artificial-bee-colony optimization with binary search equations and a neural network, evaluated on NSL-KDD for IoT intrusion detection, reporting improved detection ability and categorization relative to compared methods [37]. Sarveshwaran et al. combined Weightiness-espoused Feature Assortment (WEFA) with a Binarized Spiking Neural Network in a blockchain-integrated framework, reporting accuracy improvements of 12.94–17.99 percentage points, computation-time reductions of 31.45–59.9%, and AUC improvements of 0.83–5.43 percentage points relative to existing compared methods [38]. Binbusayyis proposed a hybrid VGG19 and 2D-CNN architecture with Radial Basis Function Support Vector Regression (RBFSVR) for dimensionality reduction, targeting fog-to-cloud SYN-flood and related attacks, evaluated across UNSW-NB15, CICIDS2017, and CICIDS2018, and reported outperformance of comparable techniques across detection rate, F-score, precision, recall, false-alarm rate, and accuracy [40]. Mohammad et al. applied data-augmentation methods across UNSW-NB15, 5G-NIDD, FLNET2023, and CIC-IDS-2017 to improve several DL architectures' performance, finding that a comparatively simple CNN-based architecture achieved highly accurate detection (up to 91% on the augmented CIC-IDS-2017 set) while more complex architectures yielded only marginal further improvement [42]. Najafli et al. proposed a Deep Reinforcement Learning (DRL) IDS for binary detection at the fog layer combined with an ensemble method for multiclass attack identification in the cloud, evaluated on CICIDS2018, reporting comparable detection performance with significantly reduced prediction time and effective handling of underrepresented attack classes [35]. Reka et al. proposed a multi-head self-attention gated graph convolutional network for multi-attack intrusion detection in MANET, implemented and evaluated in the NS-2 simulator, reporting percentage-point improvements in accuracy, precision, recall, and ROC-AUC relative to three named comparator methods (ESOC-MA-ID-MANET, EOS-IDS, IHO-MA-ID-MANET) [40].

Comparative findings: this is the largest and most heterogeneous category in the reviewed sample. Reported accuracies cluster tightly near the top of the observed range (98.75–99.995%) for studies using recent cloud/IoT-native datasets (CICIoT2023, TON_IoT), while efficiency gains — training/testing-time reduction — are most dramatic in the metaheuristic-optimization studies (Sarkar et al.; Sarveshwaran et

al.), suggesting that the main marginal value of hybridization in this literature is now computational efficiency rather than further accuracy gains, which are already near-saturated on the commonly-used benchmarks.

Strengths: metaheuristic feature selection substantially reduces both dimensionality and training/testing time, directly addressing the scalability challenge central to cloud deployment; hybrid CNN-LSTM and attention-gated GNN architectures capture both spatial and temporal/relational attack structure that single-architecture models miss.

Weaknesses: near-saturated accuracy on standard benchmarks (several studies exceed 99%) raises the concern that these datasets may no longer meaningfully discriminate between competing hybrid architectures, and may instead reflect dataset-specific overfitting rather than generalizable cloud-threat detection capability; comparator baselines vary study-to-study (some compare against ten external ML/DL baselines, others against three internally-named comparator methods), which weakens cross-study synthesis.

Unresolved gap: none of the eight hybrid/metaheuristic studies reviewed here evaluates cross-dataset generalization (training on one cloud/IoT dataset and testing on another), despite this being the most direct test of whether near-100% reported accuracy reflects genuine detection capability or benchmark-specific overfitting.

5.5 Federated-learning cloud IDS

Mechanism: federated learning trains a shared detection model across distributed cloud tenants, edge nodes, or IoT-cloud gateways without centralizing raw traffic data, directly targeting the multi-tenant data-privacy constraint identified in Section 1.

The Ioannou et al. MIoT pipeline discussed in Section 5.1 is, to this review's evidence, the clearest fully-worked federated-learning cloud-IDS example in the verified core set, using FL specifically to keep a remote cloud model current with distributed edge-classifier updates while preserving tenant data locality [41]. The broader reference base assembled for this review (Section 1.3.4, tier 2) includes additional 2024–2025 federated-IDS titles addressing IoV, medical-IoT, and hypervisor-level federated anomaly detection, which this review references narratively in Section 9 as evidence of a growing but still comparatively small research stream, without re-verifying their specific quantitative claims here.

Strengths: directly addresses the privacy/multi-tenancy constraint that is structurally unique to cloud (as opposed to generic network) IDS; naturally compatible with cross-provider or cross-tenant collaborative defense scenarios.

Weaknesses: federated aggregation introduces new attack surface (model-poisoning by a malicious participant) that is rarely evaluated alongside detection accuracy in the reviewed literature; communication overhead of federated rounds is seldom benchmarked against the centralized alternative it replaces.

Unresolved gap: poisoning-robustness of federated cloud-IDS aggregation is, in the verified core set, unaddressed — this is taken up as a roadmap pillar in Section 9 and connects directly to the general-cybersecurity poisoning-attack literature on AI/XAI systems [16].

5.6 Graph-based and transformer/attention-based cloud IDS

Mechanism: GNNs represent traffic or system relationships as graphs (hosts/flows as nodes, connections as edges) to capture relational attack structure such as lateral movement or botnet coordination; transformer/attention architectures apply self-attention over packet or event sequences.

Within the verified core set, Reka et al.'s multi-head self-attention gated graph convolutional network (Section 5.4) is the clearest example combining both paradigms, applied to MANET rather than a cloud-provider environment specifically [40]. Beyond this, this review did not identify a verified, quantitatively-detailed, cloud-provider-specific GNN or transformer study within the core evidence set — a finding that is itself significant given how established both paradigms are in the adjacent general network-IDS and cybersecurity-XAI literature [13]. This absence is treated as a primary research gap in Section 8 and a roadmap pillar in Section 9, rather than papered over with unverified claims.

6. Comparative Study Table

Table 4 consolidates the verified core set of cloud-specific AI-IDS studies reviewed in Section 5, for which dataset, technique, and quantitative results were confirmed against the source material.

Table 4. Comparative study table (verified core set).

Study	AI category	Dataset(s)	Best reported result	Cloud-specific
Vibhute & Nakum [27]	CNN + RF feature selection	CSE-CIC-IDS2018	97.07% acc., F1 97.52%	Yes
Joraviya et al. [28]	CNN (image-encoded system calls)	Leipzig IDS 2019 (containerized)	Outperforms prior pretrained-CNN baselines (no single headline % reported)	Yes (containers)
Jansi Sophia Mary & Mahalakshmi [29]	CNN + HITL	Not specified in source	5-class detection, outperforms compared literature	Yes
Bai & Sudha [30]	PCA+SVD + DL classifier	UNSW-NB15	Outperforms compared approaches (cloud forensics)	Yes
Sarkar, Keserwani & Govil [32]	ISSA (metaheuristic) + MDBN	UNSW-NB15	99.8% acc., FAR 0.02%, training 103.51→0.108s	Yes
Najafli et al. [35]	DRL (binary) + ensemble (multiclass)	CICIDS2018 (fog-to-cloud)	Comparable detection, reduced prediction time	Yes
Hasimi et al. [36]	Feed-forward ANN	Kaggle dataset	Not headline-quantified in source	Yes
Laassar et al. [37]	BABCN (bee-colony + NN)	NSL-KDD	Improved detection/categorization vs. baselines	IoT/cloud

Study	AI category	Dataset(s)	Best reported result	Cloud-specific
Sarveshwaran et al. [38]	WEFA + Binarized Spiking NN + blockchain	Not specified in source	+12.94–17.99 pp acc., –31.45–59.9% time, +0.83–5.43 pp AUC	Yes
Binbusayyis [40]	RBFSVR + VGG19 + 2D-CNN	UNSW-NB15, CICIDS2017, CICIDS2018	Outperforms compared methods across DR/F-score/precision/recall/FAR/acc.	Fog-to-cloud
Reka et al. [40]	Multi-head attention GNN	NS-2 simulated MANET	+4.3–10.4 pp acc. vs. 3 named comparators	MANET (cloud-adjacent)
Mohammad et al. [42]	CNN + data augmentation	UNSW-NB15, 5G-NIDD, FLNET2023, CIC-IDS-2017	Up to 91% acc. (augmented CIC-IDS-2017)	Yes
Ioannou et al. [41]	Ensemble RF + One-Class SVM + Federated Learning	MIoT / Raspberry Pi testbed	99.98% (RF stage), 99.7% (SVM stage)	Yes (federated)
Kavitha et al. [43]	FEFS + RNN-TDO	Not specified in source	Outperforms RNN, DNN, RNN-GA baselines	Yes
Attou et al. [44]	RF feature selection + RBFNN	Bot-IoT, NSL-KDD	>92% acc., MCC 28%→93%	Yes
Yaras & Dener [46]	Hybrid 1D-CNN + LSTM	CICIoT2023, TON_IoT	99.995% acc. (binary, CICIoT2023)	Yes

Reading note: "Best reported result" reproduces the headline metric(s) as stated in the source material for each study; metrics are not standardized across studies (some report accuracy only, others report FAR/MCC/time), which is itself the central comparability problem discussed in Section 8.

7. Attacks and Threats on Cloud AI-IDS

A structurally distinct concern from detection performance is the robustness of the AI-IDS model itself once deployed. The general cybersecurity-XAI literature has documented four attack families against AI and XAI-based security systems — poisoning attacks (PA), evasion attacks (EA), model-extraction attacks (MEA), and membership-inference attacks (MIA) [16] — but this review's evidence base contains almost no cloud-IDS-specific instances of any of the four, despite the cloud environment introducing attack-surface features (multi-tenant training data, exposed inference APIs, federated aggregation) that make these attacks arguably more relevant, not less.

7.1 Poisoning and evasion risk in multi-tenant training

Federated and multi-tenant cloud-IDS training pipelines (Section 5.5) aggregate model updates or gradients from multiple, mutually-untrusted tenants or edge nodes. A malicious tenant can, in principle, contribute crafted updates to bias the shared model — the cloud-specific analogue of the backdoor and gradient-manipulation poisoning attacks documented against general DNN and XAI systems [16]. None of the federated studies in this review's core set evaluates this threat model.

7.2 Model-extraction and inference risk on exposed cloud IDS endpoints

Cloud IDS models are frequently exposed as internal or partner-facing inference APIs (e.g., a managed-security-service offering). Query-based model-extraction and membership-inference attacks, well-documented against gradient-based and surrogate-based XAI-enabled classifiers in the general cybersecurity literature [16], have a direct cloud analogue: an attacker with API-level query access to a cloud-hosted IDS could, in principle, reconstruct decision boundaries or infer training-set membership without ever touching the underlying infrastructure. This risk is not evaluated in any cloud-IDS study identified in this review's core set.

This section is deliberately short relative to Sections 5–6: its purpose is to document an evidence gap, not to synthesize findings that do not yet exist in the cloud-specific literature. It is expanded into a dedicated roadmap pillar in Section 9.

8. Challenges and Research Gap Matrix

Table 5 consolidates the challenges surfaced throughout Sections 5–7 into a structured gap matrix, mapping each open problem to the technique category it most affects and a corresponding proposed research direction (elaborated in Section 9).

Table 5. Research gap matrix.

Gap	Evidence	Affected technique category	Proposed direction
No standardized cross-dataset evaluation	Every reviewed study evaluates on 1–4 datasets in isolation (Section 5, Table 4); no cross-dataset generalization test found	All	Mandatory cross-dataset benchmarking protocol (Section 9.5)
Near-saturated benchmark accuracy obscures real capability	Multiple hybrid-DL studies exceed 99% on standard benchmarks (Section 5.4)	Hybrid DL, CNN, ensemble	Adversarial/out-of-distribution stress-testing (Section 9.5)
Explainability essentially absent	Zero reviewed cloud-IDS studies report SHAP/LIME/Grad-CAM or comparable XAI output, despite maturity of these methods in general cybersecurity-XAI literature [13]	All, esp. CNN, ensemble	XAI-integrated cloud IDS (Section 9.1)
Federated aggregation poisoning risk unaddressed	No reviewed federated study (Section 5.5) evaluates malicious-tenant poisoning	Federated learning	Robust federated IDS (Section 9.2)
Graph/relational and transformer methods underrepresented for cloud specifically	Only one core-set study combines attention + graph structure, and not in a cloud-provider setting (Section 5.6)	GNN, transformer	GNN-based and transformer/LLM-assisted cloud IDS (Sections 9.3–9.4)
No concept-drift evaluation over time	No reviewed recurrent-architecture study (Section 5.3) tests performance degradation as traffic distributions shift	RNN/LSTM/GRU, all sequential models	Concept-drift-aware IDS (Section 9.5)

Gap	Evidence	Affected technique category	Proposed direction
Model-extraction/membership-inference risk on cloud-exposed IDS APIs unaddressed	No reviewed study evaluates query-based extraction/inference attacks against a deployed cloud IDS (Section 7.2)	All, esp. federated and API-exposed models	Adversarial-robustness-by-design for cloud IDS (Section 9.5)

9. Future Research Roadmap

Building on the gap matrix above, this section proposes five research pillars for the next generation of cloud-native AI-IDS.

9.1 XAI-integrated cloud IDS

Current state: general cybersecurity XAI (SHAP, LIME, Grad-CAM, attention-based explanation) is well-developed for malware and generic network-intrusion detection [13], but essentially unreported in the cloud-IDS-specific literature reviewed here.

Open problem: cloud SOC analysts need actionable, tenant-safe explanations (e.g., which flow features triggered an alert) without leaking cross-tenant information through the explanation channel itself — a constraint that does not exist in single-tenant network IDS.

Proposed direction: adapt post-hoc, model-agnostic techniques (SHAP/LIME) to cloud-IDS pipelines with an explicit tenant-isolation constraint on explanation outputs, and benchmark CNN-specific techniques (Grad-CAM) against the image-encoded cloud-IDS representations already in use (Section 5.2).

9.2 Federated IDS

Current state: federated learning has one clearly-worked cloud-IDS example in this review's core set (Ioannou et al., Section 5.5), plus a growing but unverified broader literature stream.

Open problem: robustness of federated aggregation to malicious or compromised tenant updates is unevaluated across the reviewed literature.

Proposed direction: robust-aggregation federated IDS (e.g., Byzantine-resilient aggregation) evaluated explicitly against a poisoning threat model, benchmarked on multi-tenant traffic simulations built from the cloud-representative datasets identified in Table 3.

9.3 Graph Neural Network (GNN)-based IDS

Current state: GNN-based detection is well-established for botnet and lateral-movement detection in the general network-IDS/XAI literature [13], but the reviewed cloud-specific literature contains only one attention-augmented graph model, applied to MANET rather than a cloud-provider environment (Section 5.6).

Open problem: cloud-native relational structures (VM-to-VM, container-to-container, service-mesh call graphs) are a natural fit for GNN modeling but remain unexplored.

Proposed direction: GNN-based detection over cloud service-mesh/container-orchestration graphs, paired with graph-specific explainability (e.g., subgraph-importance methods) to address the XAI gap in Section 9.1 simultaneously.

9.4 Transformer and LLM-assisted detection and SOC integration

Current state: transformer/attention mechanisms appear in the reviewed literature only as a component within a hybrid GNN (Section 5.6); no reviewed study uses a large language model (LLM) for cloud-IDS classification or SOC-analyst assistance.

Open problem: cloud SOC teams face high alert volume and need natural-language triage support; the general cybersecurity-XAI literature has begun exploring LLMs as both intrusion-detectors and explanation interfaces [13], but this has not been adapted to cloud-specific telemetry (VM logs, container events, service-mesh traces).

Proposed direction: LLM-assisted SOC triage layered on top of existing cloud-IDS detection output — using the LLM as an explanation/summarization interface rather than the primary classifier, to avoid the reliability and hallucination risks of LLM-only detection — evaluated against analyst time-to-decision as a primary outcome metric.

9.5 Adversarial robustness and concept-drift adaptation

Current state: as documented in Section 7, no reviewed cloud-IDS study evaluates adversarial evasion, poisoning, model-extraction, or membership-inference robustness; no reviewed study evaluates concept-drift degradation over time (Section 5.3).

Open problem: cloud traffic distributions shift continuously (new services, changing tenant mix, evolving attacker tooling), and cloud-hosted IDS models are increasingly exposed via API, making both drift-adaptation and adversarial-robustness first-class deployment requirements rather than research afterthoughts.

Proposed direction: establish a shared cross-dataset, adversarial, and drift-aware evaluation protocol for cloud-IDS research — extending the dataset benchmark matrix in Table 3 with explicit drift and adversarial-perturbation test splits — so that future studies can be compared on robustness, not accuracy alone.

10. Conclusion

This review has examined AI-driven intrusion detection for cloud computing across five research questions. On RQ1 (technique evolution), the reviewed literature shows a clear shift from classical ML toward hybrid deep-learning and metaheuristic-optimized architectures, with federated learning, GNNs, and transformers appearing only sparsely despite their maturity elsewhere in cybersecurity. On RQ2

(dataset adequacy), cloud-representativeness is uneven: legacy datasets (KDD99, NSL-KDD) remain in use alongside genuinely cloud/IoT-native benchmarks (CICIoT2023, TON_IoT, Bot-IoT), with no shared standard for what counts as an adequately cloud-representative evaluation. On RQ3 (performance trade-offs), reported accuracies cluster near saturation (many above 99%) on standard benchmarks, while metaheuristic-optimized hybrids deliver the most dramatic computational-efficiency gains — but cross-study comparability is undermined by inconsistent metric reporting (Table 4). On RQ4 (explainability), the reviewed cloud-IDS literature reports essentially no XAI integration, a striking gap relative to the maturity of SHAP/LIME/Grad-CAM in the general cybersecurity-XAI literature. On RQ5 (robustness/adaptability), adversarial evasion, poisoning, model-extraction, membership-inference, and concept-drift risks are all structurally relevant to cloud-hosted, multi-tenant, API-exposed IDS, yet none is evaluated in the reviewed core evidence set.

These findings motivate the five-pillar roadmap in Section 9 — XAI-integrated cloud IDS, robust federated IDS, GNN-based cloud IDS, transformer/LLM-assisted detection and SOC integration, and adversarial-robustness/concept-drift-aware evaluation — as the most consequential open directions for AI-driven cloud intrusion detection research through 2026 and beyond. As cloud infrastructure continues to underpin critical digital services, closing these gaps is a prerequisite not merely for higher benchmark accuracy, but for AI-IDS that cloud operators and regulators can trust, audit, and safely deploy at scale.

Declarations

Availability of Data and Materials

This review paper's availability of data and materials for the discussed intrusion detection approaches is supported by the use of standard, publicly documented datasets (Table 3) and the source publications cited throughout. No new primary dataset was generated by this review.

Funding

This work was conducted without internal or external funding or financial support.

Declaration of Competing Interest

The authors declare no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

We express our gratitude to our respective institutions for providing the support and infrastructure required to undertake this research, and to the developers and maintainers of the publicly available datasets used throughout this review.

References

- [1] Ali, M. A., & Singh, A. (2026). Intelligent threat detection: Adaptive weighted security index and enhanced Bi-LSTM residual network for cloud protection. *International Journal of Computer Networks and Applications*, 13(2), 338–355. DOI: 10.22247/ijcna/2026/18.
- [2] Elmasry, W., Akbulut, A., & Zaim, A. H. (2021). A design of an integrated cloud-based intrusion detection system with third party cloud service. *Open Computer Science*, 11(1), 365-379.
- [3] Otoum, Y., & Nayak, A. (2021). AS-IDS: Anomaly and signature based IDS for the Internet of Things. *Journal of Network and Systems Management*, 29(3), 23.
- [4] Martins, I., Resende, J. S., Sousa, P. R., Silva, S., Antunes, L., & Gama, J. (2022). Host-based IDS: A review and open issues of an anomaly detection system in IoT. *Future Generation Computer Systems*, 133, 95-113.
- [5] Thilagam, T., & Aruna, R. (2021). Intrusion detection for network based cloud computing by custom RC-NN and optimization. *ICT Express*, 7(4), 512-520.
- [6] Rana, P., Batra, I., Malik, A., Imoize, A. L., Kim, Y., Pani, S. K., Goyal, N., Kumar, A., & Rho, S. (2022). Intrusion Detection Systems in Cloud Computing Paradigm: Analysis and Overview. *Complexity*, 2022, 3999039.
- [7] Nkongolo, M., Van Deventer, J. P., Kasongo, S. M., Zahra, S. R., & Kipongo, J. (2022). A cloud based optimization method for zero-day threats detection using genetic algorithm and ensemble learning. *Electronics*, 11(11), 1749.
- [8] Singh, P., & Ranga, V. (2021). Attack and intrusion detection in cloud computing using an ensemble learning approach. *International Journal of Information Technology*, 13(2), 565-571.
- [9] Nafea, R. A., & Almaiah, M. A. (2021). Cyber Security Threats in Cloud: Literature Review. 2021 International Conference on Information Technology (ICIT), Amman, Jordan, 779-786.
- [10] Mayuranathan, M., Murugan, M., & Dhanakoti, V. (2021). Best features based intrusion detection system by RBM model for detecting DDoS in cloud environment. *Journal of Ambient Intelligence and Humanized Computing*, 12(3), 3609-3619.
- [11] Aslan, Ö., Ozkan-Okay, M., & Gupta, D. (2021). Intelligent behavior-based malware detection system on cloud computing environment. *IEEE Access*, 9, 83252-83271.
- [12] Kumar, P., Gupta, G. P., & Tripathi, R. (2021). Design of anomaly-based intrusion detection system using fog computing for IoT network. *Automatic Control and Computer Sciences*, 55(2), 137-147.
- [13] Einy, S., Oz, C., & Navaei, Y. D. (2021). The anomaly- and signature-based IDS for network security using hybrid inference systems. *Mathematical Problems in Engineering*, 2021, 1-10.
- [14] Bhusal, D., et al. (2023). SoK: Modeling explainability in security analytics for interpretability, trustworthiness, and usability. *Proceedings of the 18th International Conference on Availability, Reliability and Security*, 1-12.
- [15] Chinnnasamy, R., Subramanian, M., Easwaramoorthy, S. V., & Cho, J. (2025). Deep learning-driven methods for network-based intrusion detection systems: A systematic review. *ICT Express*, 11(1), 181-215.
- [16] Alsoufi, M. A., et al. (2021). Anomaly-based intrusion detection systems in IoT using deep learning: A systematic literature review. *Applied Sciences*, 11(18), 8383.
- [17] Tavallaee, M., Bagheri, E., Lu, W., & Ghorbani, A. A. (2009). A detailed analysis of the KDD CUP 99 data set. 2009 IEEE Symposium on Computational Intelligence for Security and Defense Applications, 1-6.
- [18] McHugh, J. (2000). Testing intrusion detection systems: A critique of the 1998 and 1999 DARPA intrusion detection system evaluations as performed by Lincoln Laboratory. *ACM Trans. Inf. Syst. Secur.*, 3, 262-294.
- [19] Nassif, A. B., Talib, M. A., Nasir, Q., Albadani, H., & Dakalbab, F. M. (2021). Machine Learning for Cloud Security: A Systematic Review. *IEEE Access*, 9, 20717-20735.
- [20] Sharafaldin, I., Lashkari, A. H., & Ghorbani, A. A. (2018). Toward generating a new intrusion detection dataset and intrusion traffic characterization. *Proceedings of ICISSP*, 108-116.
- [21] Bhuyan, M. H., Bhattacharyya, D. K., & Kalita, J. K. (2015). Towards generating real-life datasets for network intrusion detection. *IJ Network Security*, 17, 683-701.
- [22] Sharafaldin, I., Gharib, A., Lashkari, A. H., & Ghorbani, A. A. (2018). Towards a reliable intrusion detection benchmark dataset. *Software Networking*, 2018, 177-200.
- [23] Moustafa, N., & Slay, J. (2015). UNSW-NB15: A comprehensive data set for network intrusion detection systems. 2015 Military Communications and Information Systems Conference (MilCIS), 1-6.

- [24] Koroniotis, N., Moustafa, N., Sitnikova, E., & Turnbull, B. (2019). Towards the development of realistic botnet dataset in the Internet of Things for network forensic analytics: Bot-IoT dataset. *Future Generation Computer Systems*, 100, 779-796.
- [25] Pa, Y. M. P., Suzuki, S., Yoshioka, K., Matsumoto, T., Kasama, T., & Rossow, C. (2015). IoT POT: Analysing the rise of IoT compromises. 9th USENIX Workshop on Offensive Technologies (WOOT 15).
- [26] Mirsky, Y., Doitshman, T., Elovici, Y., & Shabtai, A. (2018). Kitsune: An ensemble of autoencoders for online network intrusion detection. *arXiv:1802.09089*.
- [27] Vibhute, A. D., & Nakum, V. (2024). Deep learning-based network anomaly detection and classification in an imbalanced cloud environment. *Procedia Computer Science*, 232, 1636-1645.
- [28] Joraviya, N., Gohil, B. N., & Rao, U. P. (2024). DL-HIDS: deep learning-based host intrusion detection system using system calls-to-image for containerized cloud environment. *The Journal of Supercomputing*, 1-29.
- [29] Jansi Sophia Mary, C., & Mahalakshmi, K. (2024). Optimizing deep learning-based intrusion detection in cloud computing environment with chaotic tunicate swarm algorithm. *Journal of Intelligent & Fuzzy Systems (Preprint)*.
- [30] Bai, V. S., & Sudha, T. (2024). Deep Learning Inspired Intelligent Framework to Ensure Effective Intrusion Detection in Cloud. *International Journal of Intelligent Systems and Applications in Engineering*, 12(14s), 441-456.
- [31] Arango-López, J., Isaza, G., Ramirez, F., Duque, N., & Montes, J. (2024). Cloud-based deep learning architecture for DDoS cyber attack prediction. *Expert Systems*, e13552.
- [32] Sarkar, N., Keserwani, P. K., & Govil, M. C. (2024). A better and fast cloud intrusion detection system using improved squirrel search algorithm and modified deep belief network. *Cluster Computing*, 27(2), 1699-1718.
- [33] [Removed during PRISMA screening — record found unrelated to cloud AI-IDS scope (biochemical enzyme-kinetics study, no security content); retained here only to document the screening decision, per Section 1.3.2.]
- [34] Najafli, S., Haghighat, A. T., & Karasfi, B. (2024). A novel reinforcement learning-based hybrid intrusion detection system on fog-to-cloud computing.
- [35] Hasimi, L., Zavanis, D., Shakshuki, E., & Yasar, A. (2024). Cloud Computing Security and Deep Learning: An ANN approach. *Procedia Computer Science*, 231, 40-47.
- [36] Laassar, I., Hadi, M. Y., Arifullah, H. R., & Khan, F. S. (2024). Proposed algorithm base optimisation plan for feature selection-based intrusion detection in cloud computing. *Indonesian Journal of Electrical Engineering and Computer Science*, 33(2), 1140-1149.
- [37] Sarveshwaran, V., Pandiaraj, S., Bindu, G., Ganesan, V., & Swamidason, I. T. (2024). Binarized Spiking Neural Network with blockchain based intrusion detection framework for enhancing privacy and security in cloud computing environment. *Applied Soft Computing*, 154, 111218.
- [38] Binbusayyis, A. (2024). Hybrid VGG19 and 2D-CNN for intrusion detection in the FOG-cloud environment. *Expert Systems with Applications*, 238, 121758.
- [39] Reka, R., Karthick, R., Ram, R. S., & Singh, G. (2024). Multi head self-attention gated graph convolutional network based multi-attack intrusion detection in MANET. *Computers & Security*, 136, 103526.
- [40] Ioannou, I., Nagaradjane, P., Angin, P., Balasubramanian, P., Kavitha, K. J., Murugan, P., & Vassiliou, V. (2024). GEMLIDS-MIOT: A Green Effective Machine Learning Intrusion Detection System based on Federated Learning for Medical IoT network security hardening. *Computer Communications*.
- [41] Mohammad, R., Saeed, F., Almazroi, A. A., Alsubaei, F. S., & Almazroi, A. A. (2024). Enhancing Intrusion Detection Systems Using a Deep Learning and Data Augmentation Approach. *Systems*, 12(3), 79.
- [42] Kavitha, C., Gadekallu, T. R., K, N., Kavin, B. P., & Lai, W. C. (2023). Filter-based ensemble feature selection and deep learning model for intrusion detection in cloud computing. *Electronics*, 12(3), 556.
- [43] Attou, H., Mohy-eddine, M., Guezaz, A., Benkirane, S., Azrour, M., Alabdultif, A., & Almusallam, N. (2023). Towards an intelligent intrusion detection system to detect malicious activities in cloud computing. *Applied Sciences*, 13(17), 9588.
- [44] Chen, S., et al. (2023). Ride: real-time intrusion detection via explainable machine learning implemented in a memristor hardware architecture. 2023 IEEE Conference on Dependable and Secure Computing (DSC), 1-8.
- [45] Yaras, S., & Dener, M. (2024). IoT-Based Intrusion Detection System Using New Hybrid Deep Learning Algorithm. *Electronics*, 13(6), 1053.