Bio-Inspired Feature Selection and Deep Learning for DDoS Detection: A Systematic Review

Shakirah Saidin[1], Syifak Izhar Hisham[*2]

[1] Faculty of Computing and Information Technology, Tunku Abdul Rahman University of Management and Technology, Kuantan, MALAYSIA

[1,2] Faculty of Computing, Universiti Malaysia Pahang Al-Sultan Abdullah, Pekan, MALAYSIA

DOI: https://dx.doi.org/10.47772/IJRISS.2026.100500633

Received: 09 May 2026; Accepted: 14 May 2026; Published: 09 June 2026

ABSTRACT

Distributed Denial of Service (DDoS) attacks represent one of the most persistent and damaging cyber threats facing cloud computing environments, with global attack volumes increasing from 7.9 million incidents in 2018 to over 15 million in 2023. The high dimensionality of modern network traffic datasets which often containing hundreds of statistical features, poses a significant challenge for building efficient and accurate intrusion detection systems (IDS). This paper presents a systematic review of bio-inspired metaheuristic algorithms for feature selection in DDoS detection, coupled with deep learning classification architectures. A structured literature search was conducted across IEEE Xplore, Scopus, Web of Science, and Google Scholar using defined inclusion and exclusion criteria, yielding a final corpus of studies published between 2015 and 2025. Unlike existing reviews that treat feature selection and deep learning classification in isolation, this review uniquely integrates both dimensions within a unified analytical framework, with particular emphasis on cloud-specific DDoS scenarios, multi-stage hybrid feature selection pipelines, and cascaded deep learning architectures operating on learned latent representations. The review examines the landscape of DDoS attack taxonomies, evaluates eight widely used benchmark datasets, and provides a comparative analysis of seven bio-inspired feature selection algorithms including Genetic Algorithm (GA), Particle Swarm Optimization (PSO), Grey Wolf Optimizer (GWO), Ant Colony Optimization (ACO), Whale Optimization Algorithm (WOA), Firefly Algorithm (FA), and Salp Swarm Algorithm (SSA). Deep learning architectures surveyed include Deep Neural Networks (DNN), Convolutional Neural Networks (CNN), Long Short-Term Memory (LSTM), Bidirectional LSTM (BiLSTM), and hybrid cascaded models. A critical analysis of reported performance metrics is provided, addressing methodological concerns including dataset imbalance, overfitting risks, and limited cross-dataset generalisation. Seven key research gaps are identified, and future research directions encompassing multi-stage hybrid feature selection pipelines, multi-class detection taxonomies, and explainable AI integration are proposed.

Keywords: DDoS Detection, Intrusion Detection System, Bio-Inspired Feature Selection, Deep Learning, Cloud Security

INTRODUCTION

Distributed Denial of Service (DDoS) attacks continue to pose one of the most significant cyber threats to modern digital infrastructure. These attacks exploit the distributed nature of the Internet to generate overwhelming volumes of malicious traffic directed at target systems, causing service unavailability for legitimate users. The proliferation of cloud computing across government, finance, healthcare, and education sectors has amplified the potential impact of DDoS attacks, as service disruptions can affect millions of users simultaneously. According to industry reports, the number of global DDoS attacks increased from approximately 7.9 million in 2018 to over 15.4 million in 2023 (Cisco, 2020).

Network intrusion detection systems (NIDS) serve as the frontline defence against DDoS attacks by monitoring network traffic flows and classifying them as benign or malicious (Khraisat et al., 2019). Traditional signature-based NIDS achieve high detection rates for known threats but fail against novel, polymorphic, or zero-day

attack variants. This limitation has motivated the adoption of machine learning (ML) and deep learning (DL) approaches capable of learning complex traffic patterns directly from data (Saidin & Hisham, 2023). However, the effectiveness of these approaches is fundamentally constrained by the quality of the feature representation. Modern network traffic datasets typically contain hundreds of statistical features, many of which are redundant, irrelevant, or noisy. High-dimensional feature spaces increase computational cost, degrade classifier generalisation, and obscure the discriminative patterns that distinguish attack traffic from benign traffic.

Feature selection methods address this dimensionality challenge by identifying compact, informative feature subsets that preserve or enhance classification accuracy while reducing model complexity. Among feature selection paradigms, bio-inspired metaheuristic algorithms have emerged as particularly effective wrapper-based methods due to their ability to explore large combinatorial search spaces without requiring gradient information. Algorithms such as Particle Swarm Optimization (PSO), Grey Wolf Optimizer (GWO), Genetic Algorithm (GA), and Ant Colony Optimization (ACO) have been increasingly applied to feature selection in IDS research, demonstrating competitive performance in terms of both detection accuracy and feature reduction ratios.

Despite considerable progress, several research gaps persist in the intersection of bio-inspired feature selection and deep learning for DDoS detection. First, the majority of existing studies employ single-stage feature selection, limiting the potential benefits of hybrid multi-stage approaches that combine complementary optimisation strategies. Second, most DDoS detection studies adopt binary classification (attack versus benign), neglecting the practical importance of identifying ambiguous or suspicious traffic patterns that may represent reconnaissance or early-stage attack escalation. Third, evaluations remain predominantly confined to legacy datasets such as KDD Cup 99 and NSL-KDD, with limited adoption of newer cloud-specific datasets. Fourth, cascaded deep learning architectures, where secondary classifiers operate on learned latent representations of primary models, remain largely unexplored for DDoS detection.

While several existing reviews have addressed intrusion detection, feature selection, or deep learning individually (Khraisat et al., 2019; Thakkar & Lohiya, 2021), this review differentiates itself through four distinctive contributions. First, it provides a unified analytical framework that integrates bio-inspired feature selection and deep learning classification within a single review scope, rather than treating these dimensions in isolation. Second, it places particular emphasis on cloud-specific DDoS detection challenges, including virtualisation overhead, multi-tenancy, and elastic traffic patterns, which are underrepresented in existing surveys. Third, it identifies and analyses multi-stage hybrid feature selection pipelines as a distinct research direction, highlighting the complementary strengths of combining population-based exploration with constrained local refinement. Fourth, it critically examines cascaded deep learning architectures that operate on learned latent representations, an emerging paradigm that remains largely unaddressed in the review literature.

This paper presents a systematic review of bio-inspired feature selection methods and deep learning architectures for DDoS detection. The objectives are fourfold: (1) to survey the DDoS attack landscape and its implications for cloud security; (2) to evaluate benchmark datasets commonly used in IDS research; (3) to provide a comparative analysis of bio-inspired feature selection algorithms; and (4) to examine deep learning architectures and identify research gaps and future directions.

REVIEW METHODOLOGY

This review follows a structured literature search and selection process to ensure comprehensive and reproducible coverage of the relevant research landscape. The search was conducted across four major academic databases: IEEE Xplore, Scopus, Web of Science, and Google Scholar. The search queries combined terms from three thematic domains: (a) attack context (“DDoS,” “denial of service,” “intrusion detection,” “network intrusion”); (b) feature selection (“feature selection,” “bio-inspired,” “metaheuristic,” “swarm intelligence,” “evolutionary algorithm”); and (c) classification (“deep learning,” “neural network,” “CNN,” “LSTM,” “DNN”). Boolean operators (AND, OR) were used to combine terms across and within domains.

The inclusion criteria required that papers: (1) were published in peer-reviewed journals or conference proceedings between 2015 and 2025; (2) addressed DDoS or network intrusion detection as a primary research objective; (3) employed at least one bio-inspired or metaheuristic feature selection method, or at least one deep

learning classification architecture, or both; and (4) were written in English. The exclusion criteria removed papers that: (1) focused exclusively on signature-based or rule-based detection without any learning component; (2) addressed application-layer security issues unrelated to DDoS or network intrusion (such as malware analysis or phishing); (3) were duplicate publications or extended abstracts without full experimental methodology; or (4) were published before 2015, except for seminal works that introduced foundational algorithms (such as PSO, GWO, and GA) which were retained as essential background references.

The initial database search returned a combined pool of candidate papers after removing duplicates. Papers were screened in two stages: first by title and abstract to assess topical relevance, and then by full-text review to verify compliance with all inclusion criteria. The final corpus of retained studies forms the basis of the analysis presented in this paper. Additionally, seminal algorithm papers such as (Holland, 1975; Kennedy & Eberhart, 1995; Mirjalili et al., 2014) and authoritative technical reports (Cisa et al., 2022) were included as foundational references regardless of publication date, given their continued relevance to the field.

DDoS Attack Taxonomy and Cloud Security Landscape

A. Classification of DDoS Attacks

DDoS attacks are broadly classified according to the network protocol layer they exploit. The Cybersecurity and Infrastructure Security Agency (CISA) identifies three primary categories: volumetric attacks, protocol attacks, and application-layer attacks (Cisa et al., 2022). Volumetric attacks, including UDP flood, ICMP flood, and DNS amplification, aim to exhaust the target's bandwidth by flooding it with massive traffic volumes. These attacks are the most common and can generate traffic exceeding several terabits per second. Protocol attacks, such as SYN flood, Smurf attack, and fragmented packet attacks, exploit weaknesses in network protocol handling to consume server resources and connection state tables (Saidin et al., 2018). Application-layer attacks, including HTTP flood, Slowloris, and RUDY attacks, target the application layer by mimicking legitimate user behaviour, making them particularly difficult to detect using traditional signature-based methods.

The evolution of DDoS attacks has been characterised by increasing sophistication, multi-vector coordination, and exploitation of emerging technologies such as Internet of Things (IoT) botnets. The Mirai botnet, which compromised over 600,000 IoT devices to launch a record-breaking attack against Dyn DNS in October 2016, demonstrated the potential for IoT-based DDoS attacks to disrupt critical Internet infrastructure at scale. More recently, cloud-based DDoS attacks exploit the shared virtualised resources and multi-tenant architectures inherent to cloud computing, introducing unique attack surfaces that require specialised detection mechanisms.

B. Cloud-Specific DDoS Challenges

Cloud computing environments introduce several distinctive challenges for DDoS detection. The elastic and multi-tenant nature of cloud infrastructure means that DDoS attacks targeting one tenant may degrade performance for co-located tenants through resource contention. The virtualisation layer introduces additional latency and complexity in traffic monitoring, as network traffic may traverse virtual switches, hypervisors, and software-defined networking (SDN) controllers before reaching analysis points. Furthermore, the legitimate traffic patterns in cloud environments are inherently dynamic due to auto-scaling, load balancing, and burst traffic from legitimate marketing events or product launches, making it more difficult to establish stable baselines for anomaly detection.

In the Malaysian context, CyberSecurity Malaysia has documented a consistent year-on-year increase in DDoS incidents targeting national digital infrastructure, particularly in the government, financial services, and education sectors. The Malaysian government's push toward digital transformation through initiatives such as MyDIGITAL underscores the critical importance of robust DDoS detection capabilities to protect cloud-hosted government services and citizen-facing applications. Recent comparative analyses of cybersecurity awareness frameworks have further highlighted the need for comprehensive security strategies that combine technical detection mechanisms with organisational awareness programmes (Saidin et al., 2025).

Benchmark Datasets for DDoS and IDS Research

The selection of appropriate benchmark datasets is crucial for the training, validation, and comparative evaluation of IDS models. Table 1 summarises the most widely used benchmark datasets in DDoS and intrusion detection research, highlighting their temporal evolution, feature dimensionality, class granularity, and domain focus.

Table 1: Summary of Benchmark Datasets for DDoS and Intrusion Detection Research

Dataset	Year	Samples	Features	Classes	Focus
KDD Cup 99	1999	4.9M+	41	5	General intrusion
NSL-KDD	2009	148,517	41	5	Refined KDD
UNSW-NB15	2015	2.5M+	49	10	Modern attacks
CICIDS2017	2017	2.8M+	80	15	Diverse network attacks
CSE-CIC-IDS2018	2018	16M+	80	7	Large-scale network
CICDDoS2019	2019	50M+	88	13	DDoS-specific
BoT-IoT	2019	73M+	46	4	IoT botnets
X-IIoTID	2022	820K+	63	19	Industrial IoT

The KDD Cup 99 dataset, despite being the most widely cited benchmark in IDS literature, suffers from well-documented limitations including redundant records, unrealistic traffic distributions, and outdated attack signatures that are no longer representative of modern threat landscapes (Thakkar & Lohiya, 2021). NSL-KDD addressed some of these limitations by removing duplicate records and rebalancing class distributions but retains the same 41-feature set and attack taxonomy from the original 1998 DARPA data collection.

The UNSW-NB15 dataset represents a significant advancement by incorporating modern attack types (fuzzers, shellcode, worms, backdoors, analysis) alongside nine additional feature categories. The CIC-family datasets, CICIDS2017, CSE-CIC-IDS2018, and CICDDoS2019, generated by the Canadian Institute for Cybersecurity, provide increasingly specialised and large-scale datasets. CICDDoS2019, in particular, is dedicated to DDoS attack scenarios and includes thirteen distinct DDoS attack subcategories, making it particularly suitable for fine-grained DDoS classification research. The BoT-IoT dataset targets IoT botnet scenarios, while the X-IIoTID dataset extends coverage to Industrial IoT environments with nineteen attack classes.

Despite the availability of diverse datasets, a notable gap exists in cloud-specific DDoS datasets that capture the unique traffic patterns, virtualisation artefacts, and multi-tenancy characteristics of cloud computing environments. Recent datasets such as the BCCC-cPacket-Cloud-DDoS series have begun to address this gap but remain underexplored in the research literature. The limited adoption of cloud-native datasets constrains the ecological validity of IDS evaluations and highlights the need for greater community engagement with these emerging benchmarks.

Feature Selection Methods for Intrusion Detection

A. Feature Selection Paradigms

Feature selection methods are broadly categorised into three paradigms: filter, wrapper, and embedded approaches (Yu & Liu, 2004). Filter methods evaluate feature relevance independently of any specific classifier, using statistical measures such as Information Gain, Chi-Square, Mutual Information, and correlation coefficients to rank and select features. These methods are computationally efficient and classifier-agnostic but may miss important feature interactions since they evaluate features individually or in pairwise combinations. Common filter methods applied in IDS research include Information Gain Ratio, Chi-Square tests, and Relief-based algorithms.

Wrapper methods evaluate candidate feature subsets using the performance of a target classifier as the fitness function, thereby capturing feature interactions and their combined effect on classification accuracy (Saidin & Hisham, 2023). The search strategy for wrapper methods ranges from exhaustive search (intractable for high-dimensional datasets) to heuristic and metaheuristic search strategies. Bio-inspired metaheuristic algorithms fall within the wrapper category and have emerged as the dominant approach for feature selection in modern IDS research due to their ability to navigate large combinatorial search spaces effectively.

Embedded methods integrate feature selection within the model training process itself. Examples include L1 (Lasso) regularisation, which drives irrelevant feature coefficients to zero during optimisation, and tree-based feature importance measures from Random Forest or Gradient Boosted Tree ensembles. Embedded methods offer a compromise between computational efficiency and interaction modelling, but their feature selection is inherently tied to the specific model architecture.

B. Bio-Inspired Metaheuristic Algorithms for Feature Selection

Bio-inspired metaheuristic algorithms model natural phenomena such as swarm behaviour, evolutionary processes, and ecological interactions to perform population-based stochastic search. Their application to feature selection in IDS has grown substantially in recent years, driven by their effectiveness in exploring high-dimensional binary search spaces without gradient information. Table 2 provides a comparative summary of the most widely applied bio-inspired algorithms in IDS feature selection.

Table 2: Comparison of Bio-Inspired Feature Selection Algorithms for IDS

Algorithm	Inspiration	Exploration	Exploitation	Key Strengths
GA	Natural selection	High	Moderate	Global search, crossover diversity; high computational cost
PSO	Bird flocking	Moderate	High	Fast convergence; susceptible to premature convergence in high dimensions
GWO	Wolf hunting	Moderate	High	Hierarchical leadership; balanced exploration-exploitation transition
ACO	Ant foraging	High	Moderate	Pheromone-based learning; effective for discrete problems; smallest subsets
WOA	Whale hunting	High	Moderate	Spiral and shrinking encircling; effective on multimodal spaces
FA	Firefly flashing	High	Moderate	Light intensity attraction; naturally handles multimodal landscapes
SSA	Salp chains	Moderate	High	Leader-follower dynamics; simple parameterisation

Genetic Algorithm (GA), inspired by Darwinian natural selection, operates through selection, crossover, and mutation operators on a population of candidate solutions (Holland, 1975). GA has been widely applied to IDS feature selection, with studies reporting accuracy improvements of 1–3% alongside feature reductions exceeding 50% on benchmark datasets. However, GA's crossover-based search introduces significant computational overhead, particularly for large population sizes and high-dimensional feature spaces.

Particle Swarm Optimization (PSO), introduced by (Kennedy & Eberhart, 1995), models the social foraging behaviour of bird flocks. Each particle in the swarm adjusts its position based on its own best-known position (cognitive component) and the swarm's global best position (social component). PSO has been extensively applied to IDS feature selection, with (Almomani, 2020) comparing PSO, GWO, Firefly Algorithm (FFA), and GA on the UNSW-NB15 dataset, finding that all four algorithms achieved comparable detection accuracy while significantly reducing feature dimensionality. However, PSO is susceptible to premature convergence in high-dimensional, rugged fitness landscapes when cognitive and social acceleration coefficients are not carefully tuned.

Grey Wolf Optimizer (GWO), proposed by (Mirjalili et al., 2014), simulates the social hierarchy and cooperative hunting behaviour of grey wolf packs. The wolf population is stratified into four ranks: alpha (best solution), beta, delta, and omega, with omega wolves updating their positions by interpolating toward the three leaders through encircling behaviour. GWO has demonstrated faster convergence and higher-quality solutions than PSO on benchmark optimisation functions and has been increasingly adopted for IDS feature selection. (Sezgin et al., 2025) conducted a multi-objective comparison of GWO, GA, PSO, and ACO on the X-IIoTID dataset, finding that GWO achieved the best accuracy–subset balance with 99.50% accuracy using only 22 features (65% reduction), outperforming PSO and GA in terms of feature compactness.

Ant Colony Optimization (ACO), inspired by ant foraging pheromone trails, has shown particular effectiveness in producing the smallest feature subsets. (Sezgin et al., 2025) reported ACO selecting only 7 features (89% reduction) on the X-IIoTID dataset, though at a moderate accuracy trade-off (97.65%). More recent bio-inspired algorithms including Whale Optimization Algorithm (WOA), Firefly Algorithm (FA), Salp Swarm Algorithm (SSA), and Harris Hawks Optimization (HHO) have been applied to IDS feature selection with promising results, though their adoption remains less mature compared to GA, PSO, and GWO.

C. Hybrid and Multi-Stage Feature Selection

A growing body of research has explored hybrid feature selection approaches that combine multiple algorithms or paradigms to leverage their complementary strengths. Filter-wrapper hybrids apply a computationally inexpensive filter method to reduce the initial feature space before applying a wrapper-based metaheuristic search on the reduced subset. Multi-stage approaches chain two or more metaheuristic algorithms sequentially, where the first stage performs broad exploration and the second stage performs constrained refinement within the first stage's validated subspace. (Bakro et al., 2024) demonstrated the effectiveness of hybrid bio-inspired feature selection with Random Forest classification for cloud IDS, achieving 99.3% accuracy on CSE-CIC-IDS2018.

Despite these advances, multi-stage approaches combining complementary metaheuristic algorithms remain relatively uncommon in the literature. The majority of studies employ a single bio-inspired algorithm in isolation, potentially missing the benefits of hierarchical search strategies that balance population-based exploration with constrained local refinement.

Deep Learning Architectures for DDoS Detection

A. Deep Neural Networks (DNN)

Deep Neural Networks with multiple fully connected hidden layers serve as the foundational deep learning architecture for network intrusion detection. DNNs learn hierarchical feature representations through successive nonlinear transformations, enabling them to capture complex relationships among network traffic features. Regularisation techniques including Dropout, Batch Normalisation, and weight decay are commonly employed

to mitigate overfitting on imbalanced IDS datasets. (Alzughaibi & El Khediri, 2023) demonstrated DNN with PSO-based feature selection achieving 98.7% accuracy on CSE-CIC-IDS2018. DNNs provide strong baseline performance and serve as the foundation for more complex architectures, though their fully connected architecture does not explicitly model spatial or temporal relationships among features.

B. Convolutional Neural Networks (CNN)

Convolutional Neural Networks have been adapted for tabular network traffic data by treating feature vectors as one-dimensional signals over which convolutional filters detect local feature correlations. One-dimensional CNNs apply learned filters of small kernel sizes (typically 3–5) across the feature vector, detecting co-activation patterns among adjacent feature dimensions. This approach is particularly effective when features are arranged in semantically meaningful groups, for example, flow statistics, timing statistics, and payload statistics occupying contiguous positions in the feature vector. (Ferrag et al., 2021) compared DNN, CNN, and RNN architectures for DDoS detection, finding that CNN outperformed traditional machine learning classifiers for binary DDoS classification on CICDDoS2019, achieving accuracy exceeding 99% through spatial correlation modelling that conventional approaches cannot exploit.

Two-dimensional CNN adaptations have also been explored, where network traffic features are reshaped into square or rectangular matrices to leverage standard image-classification architectures. However, this reshaping is often arbitrary and may not reflect meaningful spatial relationships, potentially introducing spurious correlations. Recent work has explored CNNs operating on learned latent representations of other models rather than raw input features, treating the latent vector as a structured signal amenable to convolutional analysis.

C. Recurrent Architectures: LSTM and BiLSTM

Long Short-Term Memory (LSTM) networks and their bidirectional variants (BiLSTM) have been applied to DDoS detection to capture sequential dependencies in network traffic flow features. LSTM's gated memory architecture (input, forget, and output gates) enables selective retention of long-range dependencies that may characterise temporal attack patterns such as slow-rate DDoS or gradual escalation behaviour. BiLSTM extends this by processing the feature sequence in both forward and backward directions, providing richer contextual information at each position.

Hybrid CapsNet + BiLSTM models have demonstrated promising results, achieving 97–99% accuracy across multiple benchmark datasets by leveraging CapsNet's spatial feature extraction with BiLSTM's sequential learning capability (Sharma, 2025). However, LSTM-based architectures introduce higher computational cost and memory requirements compared to CNNs, and their sequential processing nature limits parallelisation during training.

D. Emerging Architectures

Several emerging deep learning architectures have been applied or proposed for DDoS detection. Deep Reinforcement Learning (DRL) frameworks model the detection problem as a sequential decision-making task, where agents learn optimal detection policies through interaction with the network environment. Recent work (Swain, 2025) has investigated actor-critic algorithms including TD3, DDPG, and A2C for real-time cloud DDoS detection, with hybrid feature selection pipelines combining Boruta, SHAP, and cross-validation stability analysis to ensure feature robustness and interpretability.

Transformer-based architectures, which have revolutionised natural language processing and computer vision, have begun to be explored for network intrusion detection. Self-attention mechanisms enable Transformers to model long-range dependencies across feature dimensions without the sequential processing constraint of RNNs, potentially offering both improved performance and training parallelism. Graph Neural Networks (GNNs) have also been proposed for network-level intrusion detection, modelling network topology and flow relationships as graph structures. However, both Transformer and GNN architectures remain in early stages of adoption for DDoS detection.

E. Comparative Summary

Table 3 summarises representative deep learning approaches for DDoS detection from recent literature, highlighting the diversity of architectures, datasets, and reported performance.

Table 3: Summary of Representative Deep Learning Approaches for DDoS Detection

Reference	Architecture	Dataset	Accuracy	Key Finding
(Ferrag et al., 2021)	DNN, CNN, RNN	CICDDoS2019	99.0%+	CNN outperformed traditional ML for binary DDoS classification
(Hnamte et al., 2024)	DNN with SDN	CICDDoS2019	99.5%	DNN effective in SDN environments for real-time detection
(AlSaleh et al., 2024)	BaysCNN	CICDDoS2019	99.66%	Bayesian CNN achieved high multi-class DDoS detection in cloud
(Bakro et al., 2024)	RF + Bio-inspired FS	CSE-CIC-IDS2018	99.3%	Hybrid bio-inspired FS with RF effective for cloud IDS
(Alzughaibi & El Khediri, 2023)	DNN + PSO	CSE-CIC-IDS2018	98.7%	PSO-based feature selection improved DNN performance
(Sezgin et al., 2025)	Multi-obj FS (GWO, GA, PSO, ACO)	X-IIoTID	97.6–99.6%	GWO best accuracy-subset balance; ACO smallest subsets

Several observations emerge from the comparative analysis. First, reported accuracy values consistently exceed 97% on benchmark datasets, suggesting that the detection problem is relatively well-solved for binary classification on standard datasets. Second, the choice of dataset significantly influences reported performance, with older datasets (KDD, NSL-KDD) generally yielding higher accuracy than more recent and complex datasets. Third, hybrid approaches combining feature selection with deep learning consistently outperform architectures without explicit feature selection, highlighting the importance of dimensionality reduction for both performance and computational efficiency. Fourth, evaluation on cloud-specific datasets remains rare, limiting confidence in the generalisability of reported results to real-world cloud environments.

F. Critical Analysis of Reported Performance Metrics

While the high accuracy values reported across the surveyed studies are encouraging, several methodological concerns warrant critical examination. First, dataset imbalance is a pervasive issue in IDS research. Many benchmark datasets exhibit significant class imbalance, with attack samples substantially outnumbered by benign traffic or vice versa. In such settings, overall accuracy can be misleading, as a classifier that overwhelmingly predicts the majority class may achieve high accuracy without detecting minority-class attacks effectively. Studies that report only accuracy without complementary metrics such as macro-averaged F1-score, per-class precision and recall, or the Matthews Correlation Coefficient (MCC) provide an incomplete picture of classifier performance, particularly for minority attack classes.

Second, overfitting risk is heightened when models are trained on high-dimensional feature spaces with limited sample diversity. Deep learning models with millions of trainable parameters can memorise training set patterns, achieving near-perfect training accuracy that does not generalise to unseen traffic distributions. The absence of rigorous cross-validation protocols, independent test sets, or k-fold evaluation in some studies raises concerns about the reproducibility and robustness of reported performance. Studies that report training accuracy without clearly separated held-out test accuracy are particularly susceptible to this limitation.

Third, data leakage represents a subtle but consequential methodological flaw. If feature scaling, normalisation, or feature selection is fitted on the entire dataset before train-test splitting, information from the test set leaks into the training process, inflating reported accuracy. Proper experimental design requires that all preprocessing steps, including standardisation, imputation, and feature selection, be fitted exclusively on the training partition. Several of the surveyed studies do not provide sufficient methodological detail to determine whether data leakage was avoided, limiting confidence in the reported results.

Fourth, the lack of cross-dataset generalisation studies is a critical gap. Models evaluated on a single dataset may exploit dataset-specific artefacts, recording conditions, or feature distributions that do not transfer to other network environments. The practical utility of an IDS model depends on its ability to generalise across heterogeneous network topologies, traffic volumes, and attack variants. Very few of the surveyed studies evaluate models on datasets different from those used during training, making it difficult to assess real-world deployment readiness. Researchers and practitioners should therefore interpret the high accuracy values reported in the literature with appropriate caution, recognising that they represent performance under controlled experimental conditions rather than guaranteed operational effectiveness.

RESEARCH GAPS AND FUTURE DIRECTIONS

Based on the systematic review presented in the preceding sections, this paper identifies seven key research gaps and corresponding future directions, summarised in Table 4.

Table 4: Identified Research Gaps and Proposed Future Directions

Research Gap	**Current Limitation**	**Future Direction**
Single-stage feature selection	Most studies use one algorithm in isolation	Multi-stage hybrid pipelines combining complementary algorithms
Binary classification dominance	Majority of studies classify only attack vs. benign	Multi-class taxonomies including suspicious or ambiguous traffic
Limited cloud-specific datasets	Most evaluations use generic network datasets	Cloud-native datasets capturing virtualisation and multi-tenancy patterns
Cascaded classifier architectures	Secondary classifiers on DNN latent representations are underexplored	CNN/LSTM/Transformer operating on learned latent features
Cross-dataset generalisation	Models evaluated on single datasets only	Transfer learning and domain adaptation across heterogeneous datasets
Computational efficiency	Bio-inspired methods have high evaluation costs	GPU-accelerated proxy evaluators and surrogate-assisted optimisation
Explainability	Deep learning models operate as black boxes	SHAP/LIME integration with feature selection for interpretable IDS

A. Multi-Stage Hybrid Feature Selection Pipelines

The predominance of single-stage feature selection in the current literature represents a significant untapped opportunity. Multi-stage pipelines that chain complementary algorithms, for example, a population-based explorer for broad dimensionality reduction followed by a hierarchical refiner for constrained local optimisation could achieve superior feature compactness without sacrificing detection performance. The key design principle is constraining the second-stage search exclusively within the first stage's validated subspace, avoiding redundant re-exploration and focusing computational effort on fine-grained refinement. Novel bio-inspired

algorithms with hierarchical population structures, such as those modelling ecological symbiotic relationships, offer promising mechanisms for multi-scale exploration at different levels of granularity.

B. Multi-Class Detection Taxonomies

The binary classification paradigm (attack versus benign) that dominates the current literature fails to capture the nuanced reality of network traffic in operational environments. A more practical detection taxonomy should include an intermediate "suspicious" or "ambiguous" class that captures borderline traffic patterns such as reconnaissance probing, low-rate flooding, or early-stage attack escalation that do not meet the threshold for definitive attack classification but warrant further investigation. Multi-class taxonomies also enable differentiated response strategies: immediate blocking for confirmed attacks, enhanced monitoring for suspicious traffic, and normal forwarding for benign traffic.

C. Cloud-Native Datasets and Evaluation

The limited adoption of cloud-specific datasets constrains the ecological validity of current IDS research. Future work should prioritise evaluation on datasets that capture cloud-specific phenomena including virtualisation overhead, multi-tenant resource contention, auto-scaling traffic bursts, and container-based microservice communication patterns. Cross-dataset generalisation studies, training on one dataset and evaluating on another, are essential for assessing the robustness and transferability of detection models across heterogeneous network environments.

D. Cascaded Classifier Architectures

The potential of secondary classifiers operating on learned latent representations of primary models remains underexplored. A DNN's penultimate hidden layer produces a compact, abstract latent representation of the input features that may contain discriminative patterns not fully exploited by the primary classifier's output layer. CNNs applied to these latent vectors can detect local co-activation patterns among adjacent latent dimensions, while LSTM variants can model sequential dependencies. Attention-based and Transformer architectures offer further possibilities for modelling global interactions within the latent space.

E. Explainability and Interpretability

As IDS are deployed in critical infrastructure environments, the ability to explain detection decisions becomes increasingly important for operator trust, regulatory compliance, and forensic analysis. The integration of model-agnostic explainability methods such as SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations) with bio-inspired feature selection pipelines offers a promising direction. SHAP-based feature importance can validate and interpret the feature selections made by metaheuristic algorithms, while LIME can provide instance-level explanations for individual detection decisions.

F. Federated and Privacy-Preserving Approaches

Federated Learning (FL) enables distributed model training across multiple network domains without sharing raw traffic data, addressing privacy concerns in multi-organisational IDS deployments. Federated bio-inspired feature selection, where each participating node runs a local feature selection process and aggregates results through a federated protocol, remains an open research direction. Privacy-preserving feature selection that maintains detection performance while protecting sensitive network topology information represents a critical requirement for real-world multi-stakeholder cloud environments.

CONCLUSION

This paper has presented a systematic review of bio-inspired feature selection methods and deep learning architectures for DDoS detection. The review examined the DDoS attack landscape and its implications for cloud computing security, evaluated eight widely used benchmark datasets, and provided comparative analyses of seven bio-inspired feature selection algorithms and multiple deep learning architectures. Key findings indicate that bio-inspired metaheuristic algorithms, particularly GWO, PSO, and GA, achieve substantial dimensionality

reduction (50–89%) while maintaining or improving detection accuracy when combined with deep learning classifiers. Among the algorithms reviewed, GWO demonstrates the strongest balance between accuracy and feature compactness, while ACO produces the smallest feature subsets at moderate accuracy trade-offs.

The review identifies seven critical research gaps: single-stage feature selection dominance, binary classification focus, limited cloud-specific dataset adoption, underexplored cascaded classifier architectures, cross-dataset generalisation deficits, computational efficiency constraints, and insufficient model explainability. Addressing these gaps through multi-stage hybrid pipelines, multi-class taxonomies, cloud-native evaluations, cascaded architectures on learned latent representations, and explainable AI integration will advance the state of the art toward more robust, efficient, and interpretable DDoS detection systems suitable for real-world cloud deployment.

The accelerating sophistication of DDoS attacks, coupled with the expanding adoption of cloud infrastructure globally and in the Malaysian context specifically, underscores the continued urgency of this research domain. Future work that integrates the identified directions, particularly multi-stage bio-inspired feature selection with cascaded deep learning classifiers evaluated on cloud-native datasets, holds significant promise for advancing practical DDoS detection capabilities.

ACKNOWLEDGEMENT

The authors would like to express our gratitude to Universiti Malaysia Pahang Al-Sultan Abdullah (UMPSA) and Tunku Abdul Rahman University of Management and Technology.

REFERENCES

1. Almomani, O. (2020). A Feature Selection Model for Network Intrusion Detection System Based on PSO, GWO, FFA and GA Algorithms. Symmetry 2020, Vol. 12, Page 1046, 12(6), 1046. https://doi.org/10.3390/SYM12061046
2. AlSaleh, I., Al-Samawi, A., & Nissirat, L. (2024). Novel Machine Learning Approach for DDoS Cloud Detection: Bayesian-Based CNN and Data Fusion Enhancements. Sensors, 24(5), 1418. https://doi.org/10.3390/s24051418
3. Alzughaibi, S., & El Khediri, S. (2023). A Cloud Intrusion Detection Systems Based on DNN Using Backpropagation and PSO on the CSE-CIC-IDS2018 Dataset. Applied Sciences, 13(4), 2276. https://doi.org/10.3390/app13042276
4. Bakro, M., Kumar, R. R., Husain, M., Ashraf, Z., Ali, A., Yaqoob, S. I., Ahmed, M. N., & Parveen, N. (2024). Building a Cloud-IDS by Hybrid Bio-Inspired Feature Selection Algorithms Along With Random Forest Model. IEEE Access, 12, 8846–8874. https://doi.org/10.1109/ACCESS.2024.3353055
5. Cisa, Fbi, & Ms-isa. (2022). Understanding and Responding to Distributed Denial of Service Attacks.
6. Cisco. (2020). Cisco Annual Internet Report - Cisco Annual Internet Report (2018–2023) White Paper - Cisco. https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html
7. Ferrag, M. A., Shu, L., Djallel, H., & Choo, K.-K. R. (2021). Deep Learning-Based Intrusion Detection for Distributed Denial of Service Attack in Agriculture 4.0. Electronics, 10(11), 1257. https://doi.org/10.3390/electronics10111257
8. Hnamte, V., Najar, A. A., Nhung-Nguyen, H., Hussain, J., & Sugali, M. N. (2024). DDoS attack detection and mitigation using deep neural network in SDN environment. Computers and Security, 138. https://doi.org/10.1016/J.COSE.2023.103661
9. Holland, J. H., & Book, A. B. (1975). ADAPTATION IN NATURAL AND ARTIFICIAL SYSTEMS.
10. Kennedy, J., & Eberhart, R. (1995). Particle swarm optimization. Proceedings of ICNN'95 - International Conference on Neural Networks, 4, 1942–1948. https://doi.org/10.1109/ICNN.1995.488968
11. Khraisat, A., Gondal, I., Vamplew, P., & Kamruzzaman, J. (2019). Survey of intrusion detection systems: techniques, datasets and challenges. Cybersecurity, 2(1). https://doi.org/10.1186/s42400-019-0038-7
12. Mirjalili, S., Mirjalili, S. M., & Lewis, A. (2014). Grey Wolf Optimizer. Advances in Engineering Software, 69, 46–61. https://doi.org/10.1016/j.advengsoft.2013.12.007

13. Saidin, S. (2018). Review of Firewall Optimization Techniques. Proceedings Book: National Conference for Postgraduate Research (NCON-PGR 2018), 28-29 August 2018 , Universiti Malaysia Pahang, Gambang, Pahang, 15–21. http://ncon-pgr.ump.edu.my/index.php/en/download/proceedings-book
14. Saidin, S. B., & Hisham, S. B. I. (2023). A Survey on Supervised Machine Learning in Intrusion Detection Systems for Internet of Things. 8th International Conference on Software Engineering and Computer Systems, ICSECS 2023, 419–423. https://doi.org/10.1109/ICSECS58457.2023.10256275
15. Saidin, S., Hisham, S. I., Jamil, A. M., & Abdullah, N. A. (2025). Comparative Analysis of Cybersecurity Awareness Frameworks: Effectiveness and Scope. 2025 IEEE International Conference on Industrial Technology & Computer Engineering (ICITCE), 28–32. https://doi.org/10.1109/ICITCE65255.2025.11210771
16. Sezgin, A., Ulaş, M., & Boyacı, A. (2025). Multi-Objective Feature Selection for Intrusion Detection Systems: A Comparative Analysis of Bio-Inspired Optimization Algorithms. Sensors 2025, Vol. 25, Page 6099, 25(19), 6099. https://doi.org/10.3390/S25196099
17. Sharma, V. (2025). Improving Intrusion Detection with Hybrid Deep Learning Models: A Study on CIC-IDS2017, UNSW-NB15, and KDD CUP 99. *Journal of Information Systems Engineering and Management*, 10(11s), 633–650. https://doi.org/10.52783/jisem.v10i11s.1665
18. Swain, P. K. (2025). Cloud-based DDoS detection using hybrid feature selection with deep reinforcement learning (DRL). *Scientific Reports*, 15(1), 36546. https://doi.org/10.1038/s41598-025-18857-3
19. Thakkar, A., & Lohiya, R. (2021). A survey on intrusion detection system: feature selection, model, performance measures, application perspective, challenges, and future research directions. In Artificial Intelligence Review (Number 0123456789). Springer Netherlands. https://doi.org/10.1007/s10462-021-10037-9
20. Yu, L., & Liu, H. (2004). Efficient Feature Selection via Analysis of Relevance and Redundancy. Journal of Machine Learning Research, 5, 1205–1224.