International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

The Role of Artificial Intelligence in Enhancing Cybersecurity Systems

Sanyam Singh

Abstract
Write this at the end; summarise AI-driven IDS, behavioural analytics, automated response, research gap, and final contribution.

Chapter 1. Introduction

Cybersecurity has entered a phase where traditional, rule-based defence mechanisms are increasingly inadequate. Modern cyber-attacks are now multi-stage, automated, and highly adaptive, leveraging complex sequences rather than single, isolated incidents (Khan, 2023). At the same time, the explosion of cloud computing, Internet of Things (IoT) devices and inter-connected networks has dramatically expanded the attack surface. As a result, organisations face an unprecedented volume and variety of threats.

Artificial Intelligence (AI) technologies have emerged as a critical tool for cybersecurity: machine learning (ML) and deep learning (DL) algorithms enable real-time anomaly detection, behavioural profiling and automated threat responses (Mohamed, 2025). Despite these advances, many security operations still rely heavily on legacy rule-based systems, which suffer from high false-positive rates, poor adaptability to new threats and limited scalability (Reddy et al., 2024).

Existing cybersecurity systems struggle to keep pace with the evolving threat landscape. Traditional signature-based and static anomaly detection solutions cannot effectively handle novel, multi-phase attacks. Consequently, there is often a lack of automation, which leads to significant dwell time for attackers inside networks. This delay increases the potential damage and loss (Cremer et al., 2022). Furthermore, behavioural anomalies and insider threats remain largely undetected by rule-based systems that do not model dynamic user behaviour or adapt over time.

The aim of this study is to analyse the role of Artificial Intelligence in enhancing cybersecurity systems, specifically by focusing on intrusion detection systems (IDS), behavioural analytics, and automated threat responses. The specific objectives are:

- To compare different DL/ML architectures and evaluate their effectiveness in IDS and behavioural analytics.
- To propose an integrated AI-driven detection–prediction–response framework that organisations can adopt.
- To identify and discuss the limitations, deployment challenges and future research directions for AI in cybersecurity.

The research questions this paper will cater to will be:

1. How do AI-driven IDS models outperform traditional systems in terms of detection accuracy, false-positive rates and adaptability?

2. What roles do behavioural analytics and user profiling play in detecting insider threats and anomalous behaviours?
3. How can AI automate cybersecurity responses in order to reduce dwell time and improve incident containment?
4. What are the limitations and challenges of using AI in cybersecurity systems (for example, data bias, adversarial attacks, explainability, scalability)?

The findings of this study are significant for multiple domains: enterprises operating hybrid/remote infrastructure, cloud service providers, organisations managing IoT ecosystems, and critical infrastructure sectors. By developing a unified AI cyber-defence architecture, this research contributes both to academic knowledge and practical applications. It aims to bridge the gap between detection technologies (IDS), behavioural analytics and automated responses, thereby creating a more robust and resilient cybersecurity posture.

Chapter 2: Literature Review

2.1 AI-Driven Intrusion Detection Systems (IDS)

Artificial Intelligence has significantly transformed Intrusion Detection Systems by replacing static rule-based approaches with dynamic learning-based models capable of identifying unseen or evolving threats (Xu et al., 2025). AI-driven IDS models fall broadly into three categories: deep learning architectures, hybrid ML-DL systems, domain-specific IDS, and specialised frameworks for IoT and critical infrastructure.

Deep learning architectures such as Convolutional Neural Networks (CNN), Recurrent Neural Networks (RNN), Long Short-Term Memory (LSTM), Gated Recurrent Units (GRU), Autoencoders, and Graph Neural Networks (GNN) have shown strong performance in intrusion detection tasks due to their ability to model complex patterns in network traffic (Awajan, 2023). CNNs are effective for spatial feature extraction, while LSTM and GRU architectures capture temporal dependencies in sequential traffic flows. Autoencoders excel in unsupervised anomaly detection by learning low-dimensional representations of normal traffic patterns. GNN-based IDS models have recently gained relevance for analysing network relationships and graph-structured cyber data (Zubairu & Rabiu, 2025).

However, these architectures are not without limitations. Deep learning models are highly dependent on labelled datasets and often suffer in real-world settings where data is noisy or imbalanced. They are also vulnerable to adversarial manipulation, where small perturbations can mislead the model into misclassification. Such weaknesses highlight the need for explainable and robust detection mechanisms (Lampe & Meng, 2023).

Hybrid IDS frameworks combine traditional machine learning classifiers; such as Support Vector Machines (SVM) and Random Forests; with deep learning feature extractors. These systems leverage DL for representation learning and ML for efficient classification, thus improving detection accuracy and reducing false positives (Awajan, 2023).

Recent research highlights the use of attention mechanisms to enhance hybrid models by allowing the system to focus on the most relevant features of network traffic (AI-Driven IDS Study, 2025). Hybrid models also address computational constraints by reducing training complexity while retaining deep learning's strength in identifying subtle malicious patterns.

Industry-specific IDS models have emerged in response to attacks on cyber-physical systems, particularly automotive networks such as the Controller Area Network (CAN) bus and automotive Ethernet. These

environments require extremely low-latency intrusion detection, as even micro-delays can impact safety-critical operations (Lampe & Meng, 2023).

Research shows that lightweight DL models can detect spoofing, replay, and manipulation attacks within in-vehicle networks. However, deployment remains complex due to hardware constraints, limited computational capacity, and the need for real-time decision-making.

Industrial IoT systems face similar issues, including strict reliability requirements and constrained processing capabilities.

IoT ecosystems introduce unique security challenges due to their scale, heterogeneity, and resource limitations. AI-based IDS models for IoT rely on lightweight ML and DL architectures designed to run on low-power devices with minimal memory (Awajan, 2023).

These models aim to detect common IoT threats such as botnets, malware injection, spoofing, and scanning attacks. Techniques such as model compression, quantisation, and edge-based inference optimisation have been widely explored to improve deployment feasibility.

The primary challenge lies in balancing detection accuracy with computational efficiency while ensuring scalability across millions of connected devices.

Critical infrastructures; particularly smart grids and Advanced Metering Infrastructure (AMI); require highly reliable intrusion detection systems. Research indicates that cyber-physical attacks in these systems often involve data tampering, false data injection, and meter manipulation (El Mrabet et al., 2019). Studies comparing supervised and unsupervised learning show that unsupervised methods are often better suited for smart grid environments because they can detect novel anomalies without requiring labelled attack data.

AI-driven IDS in this domain focus on detecting abnormal consumption patterns, tampered meter readings, and anomalous communication flows between smart devices.

Despite the advancements, deploying AI-based IDS in real-world environments remains challenging

1. Encrypted traffic inspection: With most modern network traffic being encrypted, traditional deep learning models struggle to extract meaningful features without decryption (Zubairu et al., 2025).
2. Concept drift: Real-world network behaviour changes over time, causing models trained on historical data to degrade in accuracy.
3. Scalability & data imbalance: Enterprise networks generate massive volumes of data, often with very few malicious samples, leading to severe imbalance issues.
4. Metadata-based detection: Researchers increasingly advocate hybrid models that rely on metadata such as packet timing, flow duration, and frequency rather than deep packet inspection.

2.2 Behavioural Analytics & User Behaviour Analysis (UBA)

UBA focuses on understanding and modelling normal user behaviour to detect anomalies indicative of insider threats or account compromise. Unlike signature-based systems, behaviour analytics detect threats based on how users act rather than known patterns of malicious activity (Ranjan & Kumar, 2022).

Behavioural modelling involves creating baselines from user activity such as login times, device usage, file interactions, and network connections. AI enhances this by assigning anomaly scores to deviations from established behavioural baselines (AI-Driven UBA Study, 2023).

Machine learning models identify patterns that are too complex or subtle for rule-based tools, enabling early detection of suspicious activity.

Insider attacks often involve authorised users misusing their privileges, making them difficult to detect using standard IDS. Machine learning techniques classify users as “legitimate” or “malicious” by

analysing behavioural indicators such as unusual access sequences or irregular work-hour logins (Ranjan & Kumar, 2022).

A known challenge in insider threat detection is the precision–recall trade-off, as the datasets are typically imbalanced with far fewer malicious instances.

Log-based detection uses machine learning to analyse system logs, API calls, access timestamps, and event sequences. Both supervised and unsupervised clustering methods are applied to group user behaviours and flag anomalies (Olabanji et al., 2024).

Sequence-based models, such as LSTMs, are particularly effective at identifying suspicious deviations in activity patterns over time.

Cloud environments introduce risks such as VM hopping, privilege escalation, API abuse, and cross-tenant behavioural anomalies. Behaviour analytics serve as a compensating control when cloud misconfigurations expose systems to threats (Olabanji et al., 2024).

Given the dynamic nature of cloud workloads, UBA is essential in detecting abnormal resource consumption, unusual API call frequencies, and unauthorised access to virtual machines.

Modern security operations centres integrate UBA with existing tools such as SIEM, antivirus software, EDR and firewall logs. This hybrid approach allows AI systems to prioritise alerts, reduce analyst fatigue, and enhance overall threat visibility (Shankari & Rethinavalli, 2024).

Moreover, hybrid UBA models outperform rule-based monitoring by providing predictive insights and identifying threats that would otherwise remain undetected.

2.3 Automated Threat Responses (AI Agents & Reinforcement Learning)

Automation in cybersecurity has become necessary to counter rapidly evolving threats. Reinforcement Learning (RL) and agentic AI frameworks allow systems to autonomously respond to attacks, thereby reducing dwell time and preventing lateral movement (Maka et al., 2025).

RL-based cyber defence systems learn optimal actions; such as blocking IPs, isolating hosts, updating firewall rules, or throttling traffic; based on feedback from the environment. The model is structured around states, actions, and rewards (Maka et al., 2025).

Continuous training enables these agents to adapt their responses as threats evolve.

Agentic AI systems operate through a perception → reasoning → action loop. They interpret network events, analyse potential threats, and autonomously select responses. Multi-agent coordination allows distributed networks to respond simultaneously, enhancing resilience in large-scale environments.

AI-driven response systems are particularly effective in mitigating multi-stage intrusions such as ransomware campaigns or lateral movement across networks. They can segment compromised zones, detect privilege escalation attempts, and halt suspicious propagation in real time.

Security Orchestration, Automation and Response (SOAR) platforms use AI to automate triage, modify security policies dynamically, and run adaptive playbooks. These self-learning capabilities enable faster remediation and reduce the workload on human analysts (Maka et al., 2025).

Chapter 3: Methodology

This study employs a secondary qualitative research design, focusing on the systematic collection, review, and synthesis of existing academic and industry literature on AI-driven cybersecurity. Since no primary data is generated, the research relies entirely on previously published sources to analyse how Artificial Intelligence enhances intrusion detection systems (IDS), behavioural analytics (UBA), and automated threat responses.

The design follows an integrative review approach, allowing the inclusion of empirical studies, theoretical frameworks, technical papers, and domain-specific cybersecurity reports. This enables a holistic comparison of AI techniques across multiple contexts without conducting new experiments or data collection.

3.2 Data Sources

As a secondary research study, this paper draws entirely from existing published material sourced from academically credible and industry-recognised platforms. Peer-reviewed journal articles and conference papers from databases such as IEEE, ACM, Springer, Elsevier, MDPI, and arXiv form the primary foundation of the research, providing empirical evidence, experimental findings, and conceptual frameworks related to AI-driven intrusion detection systems, behavioural analytics, and automated threat responses. To complement these sources, systematic reviews and survey papers were examined to capture broader trends, methodological patterns, and comparative insights across multiple AI approaches. Although no primary datasets were used, literature referencing benchmark datasets; including NSL-KDD, CIC-IDS 2017, UNSW-NB15, ToN-IoT, and automotive CAN bus datasets; was reviewed to understand how researchers evaluate model performance in different cybersecurity contexts. Additionally, industry whitepapers, cybersecurity guidelines, and technical frameworks issued by NIST, ENISA, MITRE ATT&CK, IBM Security, and Cisco Talos were included to incorporate practical perspectives and real-world threat intelligence. Foundational texts on artificial intelligence, machine learning, and cyber-physical systems further contributed to the theoretical grounding of the study, ensuring a comprehensive and robust synthesis of existing knowledge.

3.3 Analytical Framework

This research applies a thematic analytical framework to systematically evaluate and synthesise insights from the reviewed literature. Thematic analysis enables the identification of recurring concepts, technological patterns, and methodological similarities across studies focused on AI-driven intrusion detection systems, behavioural analytics, and reinforcement learning-based automated response mechanisms. Through this approach, the literature was organised around central themes such as detection accuracy, false-positive reduction, adaptability to evolving threats, performance in specialised environments like IoT and cloud systems, and the operational feasibility of deploying AI models at scale. Rather than comparing raw numerical results; which vary significantly across datasets and experimental setups; the analysis focuses on conceptual and qualitative comparisons, assessing how different models perform under similar problem conditions. Cross-study interpretation was guided by metrics commonly reported in the literature, including detection rates, anomaly classification performance, response speed, and model robustness. This method allows a coherent understanding of how AI contributes to cybersecurity enhancements across multiple research domains without requiring primary experimental data.

3.4 Evaluation Criteria

The evaluation of findings in this study is grounded in four major criteria identified across the reviewed literature: model performance, scalability, explainability, and operational feasibility. Model performance is assessed through how studies report accuracy, precision, recall, false-positive rates, and the ability to detect both known and novel threats. Scalability considers whether AI techniques discussed in the literature can maintain efficiency when deployed across large enterprise networks, distributed cloud platforms, or resource-constrained IoT ecosystems. Explainability evaluates the extent to which models provide interpretable outputs, a critical consideration in cybersecurity where opaque or non-transparent

decisions can lead to operational risk or mistrust among analysts. Operational feasibility examines practical deployment factors, including computational requirements, latency tolerance, integration with existing security tools such as SIEM and SOAR platforms, and the overall maintainability of AI-driven systems in real-world environments. By applying these criteria, the study ensures that the synthesized findings are not only academically grounded but also relevant to practical cybersecurity implementation.

Chapter 4. Analysis & Discussion

4.1 Comparison Across IDS Approaches

The analysis of AI-driven Intrusion Detection Systems reveals clear performance distinctions between traditional machine learning (ML), deep learning (DL), and hybrid models. ML-based systems, such as Support Vector Machines and Random Forest classifiers, have historically demonstrated strong performance on structured and well-labelled datasets, particularly when traffic features are clearly defined. However, their effectiveness drops when handling high-dimensional or evolving network traffic, as they rely heavily on manual feature engineering. In contrast, DL models such as Convolutional Neural Networks (CNN), Long Short-Term Memory networks (LSTM), Gated Recurrent Units (GRU), and Graph Neural Networks (GNN) outperform ML systems by autonomously learning complex spatial, temporal, and relational patterns within traffic flows. Among these, LSTM and GRU architectures excel at modelling sequential traffic behaviour, while GNNs show strong potential for representing network topologies and detecting relational anomalies across nodes.

Hybrid models that integrate ML classifiers with DL feature extractors show some of the most promising results, particularly in reducing false-positive rates and improving generalisation across diverse environments. Nevertheless, traditional IDS approaches still hold advantages in specific scenarios. Rule-based and signature-based IDS outperform AI-driven methods when dealing with known threats requiring precise pattern matching, especially in environments where stability and low computational overhead are priorities. These findings indicate that while AI-driven IDS significantly enhance detection capabilities, traditional systems remain valuable components of layered security strategies.

4.2 Effectiveness of Behavioural Analytics

Behavioural analytics and User Behaviour Analysis (UBA) offer substantial improvements in detecting insider threats and sophisticated attacks that bypass perimeter defences. UBA models excel because they analyse behavioural baselines ; such as typical login patterns, resource access sequences, or device usage tendencies ; and detect deviations rather than relying on predefined attack signatures. This makes them particularly effective for identifying privilege misuse, lateral movement, and accounts compromised through phishing or credential theft.

In cloud environments, behavioural analytics demonstrate even greater value. Cloud-specific anomalies ; such as unusual API call frequencies, sudden changes in virtual machine behaviour, or attempts to access cross-tenant resources ; are more effectively captured through behavioural modelling than through traditional IDS tools, which often struggle with the dynamic and elastic nature of cloud systems. UBA complements IDS at runtime by providing contextual understanding of user intent and activity flow. For example, while an IDS may detect suspicious traffic spikes, UBA identifies whether the actions correspond to legitimate user behaviour or abnormal deviations, enabling more accurate decisions and reducing false alarms. Overall, UBA enhances visibility into human-centric attack vectors that are often invisible to network-only detection methods.

4.3 Efficiency of Automated Responses

Automated threat response systems, particularly those powered by Reinforcement Learning (RL), demonstrate significant efficiency improvements compared with manual Security Operations Center (SOC) interventions. RL-based systems continuously learn optimal actions ; such as isolating compromised endpoints, modifying firewall rules, terminating suspicious processes, or throttling suspicious network flows ; based on real-time environmental feedback. This learning-driven adaptability allows RL agents to respond faster and more precisely than manual analysts, who often face alert fatigue and time constraints.

One of the most notable advantages of automated response systems is the **reduction in attacker dwell time**, which is one of the strongest predictors of breach severity. By autonomously initiating containment or mitigation actions, AI-driven systems can stop the progression of multi-stage attacks before they escalate. Real-world applications demonstrate their capability: in ransomware scenarios, automated systems can detect anomalous file encryption patterns and immediately isolate affected machines, preventing spread across the network. Similarly, in cloud environments, RL-driven policies can detect abnormal instance behaviour and automatically revoke privileges or reconfigure security groups to block lateral movement. These examples highlight the operational value of automation in modern, high-speed attack scenarios.

4.4 Towards an Integrated AI Cybersecurity Architecture

The findings from this study indicate the growing need for a unified AI-driven cybersecurity architecture that integrates IDS, UBA, and automated response mechanisms into a cohesive, layered defence model. Such an architecture would combine the strengths of each domain: IDS provides network-level threat detection, UBA delivers behavioural visibility and insider threat identification, and automated response frameworks ensure rapid containment and mitigation. Integrating these systems would enable end-to-end detection, prediction, and automated action within a single workflow.

The benefits of an integrated model include higher detection accuracy, reduced response times, continuous adaptive learning, and improved resilience against both known and unknown threats. By enabling autonomous defence, organisations can significantly reduce reliance on manual analysis and overcome resource limitations within SOC teams. However, this integration is not without risks. Increased automation introduces the potential for cascading failures, especially if false positives trigger unintended actions. AI-driven systems may also be vulnerable to adversarial attacks designed to manipulate model behaviour or degrade performance. Ensuring transparency, robustness, and human oversight is therefore critical to safe implementation.

In conclusion, the integration of IDS, UBA, and automated response represents a transformative step toward intelligent, proactive cybersecurity. While challenges remain, the convergence of these technologies offers a powerful path toward building resilient digital infrastructures capable of defending against the growing complexity of modern cyber threats.

Chapter 5. Findings

The analysis of existing literature demonstrates that Artificial Intelligence significantly enhances cybersecurity mechanisms, particularly in intrusion detection, behavioural analytics, and automated threat response. Deep learning architectures such as LSTM, GRU, CNN, and GNN consistently outperform traditional machine learning models by capturing complex temporal and relational patterns in network traffic (Awajan, 2023; Xu et al., 2025). Hybrid IDS approaches; combining ML classifiers with DL feature

extractors; further improve accuracy and substantially reduce false positives, aligning with findings reported in AI-driven IDS systematic reviews (Zubairu & Rabiu, 2025). Despite this, traditional signature-based IDS still outperform AI-based models in detecting well-known threats with fixed patterns, especially in low-resource systems where computational efficiency is critical (Lampe & Meng, 2023).

Findings also confirm the growing importance of behavioural analytics in identifying insider threats and advanced persistent anomalies. UBA techniques effectively detect unusual login activity, deviation from user-specific behavioural baselines, and privilege misuse, providing capabilities that traditional IDS alone cannot achieve (Ranjan & Kumar, 2022). In cloud environments, behavioural analytics show particular strength in identifying anomalies such as API abuse and VM hopping, supporting conclusions drawn by recent cloud-security behavioural studies (Olabanji et al., 2024). Importantly, UBA systems complement IDS by adding contextual understanding of user actions in real time.

Finally, automated response systems powered by reinforcement learning demonstrate significant operational benefits. RL-based agents continuously learn optimal containment actions; isolating compromised hosts, blocking malicious IPs, or reconfiguring firewall rules; leading to substantially reduced attacker dwell time (Maka et al., 2025). Case studies highlight the effectiveness of automated systems in mitigating ransomware propagation and cloud-based lateral movement, validating claims in recent AI-automation research (Shankari & Rethinavalli, 2024). Overall, the findings support an integrated model that combines IDS, behavioural analytics, and automated responses into a modern, adaptive defence architecture.

Chapter 6. Limitations

Despite strong results reported across the literature, several limitations constrain the effectiveness of current AI-driven cybersecurity systems. A core limitation is the dependency of deep learning models on high-quality training datasets. Many datasets commonly used; such as NSL-KDD and CIC-IDS 2017; have been criticised for outdated traffic patterns, class imbalance, and limited representation of modern attack techniques (Xu et al., 2025; Zubairu & Rabiu, 2025). As a result, model performance in controlled experiments may not generalise well to real-world enterprise networks. In addition, deep learning models remain susceptible to adversarial manipulation, where subtle traffic modifications can lead to incorrect predictions, posing a significant security risk (Lampe & Meng, 2023).

Scalability poses another constraint. DL models often require intensive computational resources, limiting their deployment in IoT and industrial environments where devices operate on minimal processing power (Awajan, 2023). Explainability is also a major concern: many AI models operate as opaque black boxes, making their decisions difficult for analysts to interpret or justify during incident investigations. This challenge is emphasised in behavioural analytics research, which acknowledges the difficulty of explaining anomaly scores derived from complex user behavioural models (Ranjan & Kumar, 2022).

Finally, this study is limited by its reliance on secondary data. The findings depend entirely on the accuracy, validity, and methodological rigor of prior studies. Variations in datasets, evaluation metrics, and experimental settings across different authors create inconsistencies that cannot be resolved within the scope of a non-experimental research design (Olabanji et al., 2024). Consequently, while the literature provides strong evidence, empirical validation through primary testing remains necessary for definitive conclusions.

Chapter 7. Future Research Directions

Future research should prioritise developing adversarially robust AI models to mitigate vulnerabilities to evasion and poisoning attacks. Studies increasingly emphasise the need for resilient IDS designs capable of maintaining performance even when confronted with adversarial inputs (Lampe & Meng, 2023; Xu et al., 2025). Another direction involves improving dataset quality. Creating updated, large-scale, and realistic datasets for IoT, cloud systems, automotive networks, and critical infrastructure would greatly enhance the reliability of AI model evaluation and cross-study comparability.

Explainability should also become a central research focus. As behavioural analytics and automated responses gain prominence, integrating Explainable AI (XAI) techniques will be essential to ensure transparency, trust, and regulatory alignment (Ranjan & Kumar, 2022). Future studies should develop interpretable behavioural models that allow analysts to understand why certain actions are flagged as anomalous.

Additionally, research should explore multi-agent reinforcement learning for distributed cybersecurity automation. Such systems could coordinate responses across large networks, improving scalability and resilience (Maka et al., 2025). Integrating IDS, UBA, and RL-driven automation into a fully unified, self-learning defence architecture also presents a promising research frontier with potential to transform organisational security operations (Shankari & Rethinavalli, 2024).

Finally, ethical, legal, and governance concerns surrounding AI-driven cybersecurity; such as privacy, accountability, and automated decision-making; require deeper exploration. Ensuring that AI systems operate safely, responsibly, and in alignment with organisational and regulatory frameworks will be crucial as adoption continues to expand (Olabanji et al., 2024).

Chapter 8. Conclusion

This study examined the growing role of Artificial Intelligence in modern cybersecurity, with a focus on intrusion detection systems (IDS), behavioural analytics, and automated threat response mechanisms. The reviewed literature clearly demonstrates that AI-driven approaches significantly enhance threat detection capability by enabling models to recognise complex spatial, temporal, and relational patterns in network traffic. Deep learning architectures such as LSTM, GRU, CNN, and GNN outperform traditional machine learning and signature-based systems in adapting to evolving attack behaviours. Similarly, hybrid models that integrate ML and DL components offer improved accuracy and reduced false-positive rates, highlighting AI's value in strengthening detection processes across diverse environments.

The analysis also emphasises the importance of behavioural modelling through User Behaviour Analytics (UBA). Unlike traditional IDS, behavioural approaches detect insider threats, credential compromise, and cloud-specific anomalies by identifying deviations from established user baselines. This human-centric perspective; focusing on how legitimate users behave rather than solely on known attack signatures; adds a critical dimension to cybersecurity. UBA thereby complements IDS by providing contextual, real-time insights that increase detection precision and reduce analytical burden on security operations teams.

In addition, automated threat response systems, particularly those driven by reinforcement learning, demonstrate the potential to transform security operations. These systems autonomously determine optimal containment actions and significantly reduce attacker dwell time, which is a key determinant of breach severity. Real-world applications indicate that automated responses can effectively mitigate ransomware propagation, prevent lateral movement, and accelerate cloud-based incident remediation.

Together, these findings present a strong case for developing integrated AI-driven cyber defence architectures that unify IDS, behavioural analytics, and automated response within a single, adaptive framework. Such integration can create a more comprehensive, predictive, and resilient defence mechanism capable of addressing both known and emerging cyber threats. The combination of network-level detection, behavioural visibility, and real-time autonomous response represents a meaningful step toward proactive cybersecurity.

However, the final analysis makes clear that while AI is a powerful enabler, it is not sufficient on its own. Effective cybersecurity still requires robust governance structures, continuous human oversight, and high-quality, up-to-date datasets to ensure accuracy, transparency, and ethical deployment. Issues such as dataset bias, model explainability, scalability, and adversarial vulnerability underscore the necessity of combining technological innovation with organisational responsibility. As cyber threats continue to evolve, the successful implementation of AI-driven security systems will depend on a balanced approach that integrates advanced algorithms with strong human, procedural, and regulatory safeguards.