MULTIDISCIPLINARY RESEARCH IN COMPUTING INFORMATION SYSTEMS

VOL 05 ISSUE 07 2025

P-ISSN: 3080-7182
E-ISSN: 3080-7190

https://mrcis.org

AI-DRIVEN INTRUSION DETECTION SYSTEMS IN CLOUD ENVIRONMENTS

Sana Mehmood [1], **Hamza Tariq** [2]

Abstract. *The rapid adoption of cloud computing has introduced new complexities in cybersecurity, particularly in detecting and mitigating intrusions in dynamic, multi-tenant environments. Traditional intrusion detection systems (IDS) rely heavily on static rule-based approaches, which are often inadequate against the evolving landscape of cyber threats. This article explores the integration of Artificial Intelligence (AI) techniques in designing intelligent, adaptive, and autonomous intrusion detection systems for cloud environments. It examines various machine learning and deep learning models used for real-time anomaly detection, data preprocessing, and decision-making automation. The study also highlights challenges such as scalability, false alarm reduction, and privacy preservation in cloud-based IDS. The findings emphasize that AI-driven IDS not only enhance detection accuracy but also enable proactive threat mitigation through continuous learning and intelligent response mechanisms.*

Keywords: *Cloud Security, Artificial Intelligence, Intrusion Detection System, Deep Learning, Anomaly Detection, Cyber Threats, Network Forensics, Machine Learning.*

INTRODUCTION

With the proliferation of cloud services, the global IT ecosystem has experienced unprecedented scalability and flexibility, but it has also become an attractive target for cybercriminals. The distributed nature of cloud infrastructures, along with shared resources and virtualization, introduces new attack vectors such as hypervisor exploitation, insider threats, and advanced persistent attacks. Conventional intrusion detection systems, primarily based on signature matching, fail to detect zero-day attacks or dynamic malicious behaviors. To overcome these limitations, researchers have turned to Artificial Intelligence (AI), which offers powerful tools to detect hidden patterns and behavioral anomalies in massive datasets. AI-driven IDS in cloud environments leverage machine learning algorithms to continuously learn from network behavior, enhance detection precision, and adapt to emerging threats in real-time. These intelligent systems

[1] *Department of Information Technology, University of the Punjab, Lahore, Pakistan.*

[2] *Department of Electrical and Computer Engineering, NUST, Islamabad, Pakistan.*

represent the next generation of cybersecurity defense mechanisms that combine scalability with cognitive computing.

The Evolution of Intrusion Detection Systems:

The evolution of Intrusion Detection Systems (IDS) can be traced through several key technological shifts that mirror the broader transformation of cybersecurity itself. In the early stages, IDS operated primarily through static, rule-based mechanisms, where predefined signatures of known attacks were compared against incoming traffic. While this approach was effective for detecting known threats, it proved inadequate against emerging, sophisticated, and polymorphic attacks that continuously evolved to bypass detection. The 1990s and early 2000s saw the introduction of statistical and heuristic methods, enabling systems to detect deviations from normal behavior patterns rather than relying solely on signatures. However, these systems still suffered from high false alarm rates and scalability limitations.

The rise of Artificial Intelligence (AI) and machine learning (ML) revolutionized IDS by enabling systems to autonomously learn from data, recognize complex attack patterns, and adapt to new threats without human intervention. AI-driven IDS utilize algorithms that perform behavioral analysis, anomaly detection, and predictive modeling to identify intrusions before they cause damage. With the advent of cloud computing, IDS architectures evolved further into distributed and collaborative models, capable of handling vast amounts of data across multiple virtual environments in real time. These systems employ deep learning, reinforcement learning, and federated learning to ensure scalability, accuracy, and data privacy.

Today, modern IDS are not just reactive tools but proactive defense mechanisms integrated with cloud orchestration platforms, capable of continuous monitoring, intelligent alert prioritization, and even automated threat response. The integration of threat intelligence feeds and blockchain-backed audit trails further enhances their transparency and reliability. As a result, IDS have transitioned from being static, standalone defense tools to dynamic, AI-powered ecosystems that form the backbone of next-generation cybersecurity in the cloud era.

Role of Machine Learning in Cloud-Based IDS:

Machine learning plays a pivotal role in enhancing the intelligence, adaptability, and efficiency of cloud-based Intrusion Detection Systems (IDS). Unlike traditional rule-based systems that rely on static signatures, ML-driven IDS models leverage data-driven learning to understand complex network behaviors and automatically distinguish between benign and malicious activities. Supervised learning algorithms such as Support Vector Machines (SVM) and Random Forests (RF) excel in pattern recognition by training on labeled datasets, enabling precise classification of known attack types. Unsupervised methods like k-Means Clustering and Autoencoders detect anomalies by identifying deviations from normal network traffic, making them particularly effective against previously unseen or zero-day attacks.

Furthermore, reinforcement learning (RL) introduces an adaptive mechanism that allows IDS to evolve through continuous feedback and experience. RL agents learn optimal detection strategies by interacting with the cloud environment, thereby improving real-time response accuracy and minimizing false positives. In distributed and multi-tenant cloud infrastructures, ML models are designed to handle heterogeneous and high-dimensional data, providing scalability and robustness.

The integration of ensemble learning—combining multiple algorithms to enhance predictive performance—further strengthens detection reliability.

Recent advancements in deep learning (DL), a subset of ML, have empowered IDS to automatically extract hierarchical features from raw network traffic without manual intervention. Models such as Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) analyze temporal and spatial correlations within cloud data streams, enabling faster and more precise intrusion detection. Moreover, federated learning allows multiple cloud nodes to collaboratively train IDS models while preserving data privacy, addressing one of the critical challenges of centralized learning systems. Overall, the application of ML in cloud-based IDS marks a transformative leap toward autonomous, scalable, and intelligent cybersecurity frameworks that adapt dynamically to the evolving cyber threat landscape.

Deep Learning for Intelligent Threat Detection:

Deep learning has revolutionized intrusion detection systems by enabling them to learn hierarchical representations of network data directly from raw inputs, thereby eliminating the need for manual feature engineering. In cloud environments, where data streams are vast, dynamic, and multidimensional, **deep neural architectures** such as Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and Long Short-Term Memory (LSTM) models excel in capturing both spatial and temporal dependencies. CNNs analyze packet structures and flow-level features, detecting local anomalies across layers of abstraction, while RNNs and LSTMs model the sequential nature of network traffic, identifying time-dependent attack behaviors such as brute-force attempts or multi-step intrusions. These models continuously adapt to new threats by learning from evolving datasets, making them ideal for **elastic and multi-tenant cloud infrastructures**.

Autoencoders and generative adversarial networks (GANs) have been employed to detect subtle deviations from normal traffic that might indicate stealthy intrusions or insider threats. Autoencoders reconstruct normal network patterns, flagging any deviations as anomalies, while GANs generate synthetic attack data to improve model robustness. In addition, **attention mechanisms** and **transformer-based models** have recently been introduced to prioritize critical features within massive cloud datasets, improving detection speed and interpretability. These techniques are particularly effective for real-time analysis of security logs, virtual machine interactions, and cross-layer correlations in distributed systems.

The scalability of DL-based IDS is enhanced by integrating them with **cloud-native technologies** such as Kubernetes and edge computing nodes, allowing distributed detection close to data sources. Combined with real-time visualization dashboards and automated response mechanisms, deep learning transforms IDS from passive monitoring systems into **proactive cyber defense frameworks** capable of predicting, identifying, and neutralizing threats before they escalate. Thus, DL not only enhances accuracy and speed but also provides **adaptive intelligence**, ensuring continuous protection in ever-evolving cloud security landscapes.

Challenges and Limitations in AI-Driven IDS:

AI-driven Intrusion Detection Systems (IDS), while transformative in enhancing cybersecurity, still encounter several technical and operational limitations that hinder their full-scale deployment in cloud environments. One of the most persistent challenges is the high rate of false positives and false negatives, which can overwhelm security teams with unnecessary alerts or, conversely, allow real threats to go undetected. This problem arises because AI models often misclassify ambiguous network behaviors due to imbalanced or incomplete training datasets. Fine-tuning these models requires continuous retraining with updated data, yet the rapid evolution of cyber threats makes maintaining dataset relevance an ongoing struggle.

Another critical challenge is data privacy and compliance. Cloud infrastructures often operate under shared environments where data belongs to multiple users or organizations. Training AI models on such multi-tenant data can violate privacy regulations like GDPR or HIPAA unless advanced techniques like federated learning or differential privacy are implemented. However, these methods can introduce trade-offs between data confidentiality and detection performance. Furthermore, computational and energy overhead poses a significant limitation, as deep learning architectures—especially CNNs and LSTMs—demand large-scale processing power and memory, which may not be feasible for real-time intrusion detection on resource-constrained nodes or edge devices.

AI-driven IDS are also vulnerable to adversarial machine learning attacks, where attackers deliberately manipulate input data or training samples to mislead the model into making incorrect classifications. Such attacks can degrade IDS performance or cause it to ignore malicious activity entirely. Additionally, the “black box” nature of many deep learning models introduces challenges in interpretability; security analysts often cannot easily explain why a particular alert was triggered. This lack of transparency undermines trust and complicates forensic analysis. To address these issues, future research must focus on developing explainable AI (XAI) frameworks, energy-efficient algorithms, and resilient training mechanisms capable of resisting adversarial manipulation. The convergence of blockchain-based auditing, privacy-preserving AI, and hybrid detection architectures promises a sustainable path toward reliable and transparent AI-driven intrusion detection in cloud ecosystems.

Future Trends and Research Directions:

The future trajectory of Intrusion Detection Systems (IDS) in cloud computing is shifting toward intelligent, decentralized, and self-evolving architectures that leverage the collective strengths of multiple emerging technologies. Hybrid frameworks integrating Deep Learning (DL), Federated Learning (FL), and Blockchain are expected to dominate next-generation IDS design. Deep learning provides the analytical depth to process complex traffic patterns and predict sophisticated multi-stage attacks, while federated learning facilitates collaborative model training across distributed cloud nodes without exposing raw or sensitive data. This decentralized training paradigm ensures user privacy and compliance with international data protection standards, a vital factor in global cloud ecosystems. Blockchain complements these technologies by creating immutable, transparent, and verifiable audit trails of all network events and IDS decisions, preventing tampering and enhancing system trustworthiness.

the adoption of Explainable AI (XAI) is becoming increasingly critical to bridge the gap between model accuracy and interpretability. By providing human-understandable reasoning for AI decisions, XAI helps cybersecurity professionals validate IDS outputs, identify biases, and comply with ethical and regulatory standards. Another promising direction is the integration of autonomous and self-learning algorithms capable of adapting to new threat patterns without manual retraining. These systems continuously refine their detection capabilities using reinforcement learning and online adaptation mechanisms.

The emergence of quantum-safe AI models, edge-enabled detection, and cross-domain knowledge sharing will redefine scalability and response times in global cloud infrastructures. Combining these technologies with Zero Trust architectures and AI-driven orchestration layers will enable proactive, predictive, and context-aware intrusion detection. Ultimately, future IDS will evolve into autonomous cyber defense ecosystems—intelligent entities capable of not just detecting but also anticipating and neutralizing threats across interconnected cloud platforms with minimal human oversight.

Integration of AI with Cloud-Native Security Architectures:

The integration of Artificial Intelligence (AI) within cloud-native security architectures represents a major advancement in designing scalable and adaptive Intrusion Detection Systems (IDS). In modern cloud infrastructures, services are increasingly deployed through microservices, containers, and serverless platforms such as Kubernetes, Docker, and OpenStack. These technologies enable flexibility and rapid deployment but also introduce new security challenges, including lateral movement of threats between containers and ephemeral attack surfaces that change dynamically. Embedding AI-driven IDS directly into these cloud-native frameworks allows real-time monitoring and automated defense mechanisms that evolve with the system.

AI models integrated with container orchestration tools like Kubernetes can detect abnormal communication patterns among microservices and identify malicious workloads based on behavioral deviations. For example, machine learning algorithms can analyze API traffic, system calls, and log data from containerized applications to detect anomalies without disrupting service continuity. Similarly, serverless computing, with its event-driven architecture, benefits from AI-based threat analytics that operate at the function level, ensuring lightweight and low-latency security enforcement.

cloud-native AI-IDS solutions employ DevSecOps principles, embedding intelligent security into the continuous integration and deployment (CI/CD) pipeline. This approach ensures that IDS components are automatically updated as new containers or services are deployed, maintaining consistent protection throughout the system's lifecycle. The use of AI-driven observability tools, such as distributed tracing and real-time telemetry analysis, enhances visibility into microservice interactions, enabling rapid anomaly detection and contextual threat correlation.

Data Preprocessing and Feature Engineering in IDS:

Data preprocessing and feature engineering are among the most crucial steps in developing reliable and accurate AI-driven Intrusion Detection Systems (IDS). Before training any machine learning or deep learning model, the raw network data collected from cloud traffic logs, firewalls, and virtual machines must undergo a **comprehensive preprocessing pipeline** to ensure data quality, consistency, and interpretability. This process typically involves **data collection**, **cleaning**, **normalization**, **feature extraction**, and **dimensionality reduction**. Since cloud environments generate massive amounts of heterogeneous data from multiple sources, preprocessing ensures that the data is standardized and free from redundancies, missing values, or irrelevant attributes that could mislead the AI model.

Data cleaning removes noise, incomplete entries, and duplicate records that might arise from network latency or packet loss. Normalization techniques—such as **min-max scaling** or **z-score standardization**—are applied to bring all features to a common scale, preventing bias toward attributes with larger numerical ranges. Once the data is normalized, **feature extraction and selection** become essential to reduce computational complexity and improve model performance. Techniques such as **Principal Component Analysis (PCA)** and **Linear Discriminant Analysis (LDA)** help in dimensionality reduction by identifying the most informative variables that capture the essence of network behavior. Additionally, **autoencoders**—a form of unsupervised deep learning—can automatically learn compressed representations of data while preserving critical patterns useful for anomaly detection.

Feature engineering also plays a vital role in identifying meaningful characteristics that differentiate normal and malicious traffic, such as packet size, connection duration, source-destination relationships, and protocol frequency. Correlation-based feature selection helps eliminate redundant variables, ensuring that only non-overlapping, high-impact features are used for model training. In modern AI-based IDS, preprocessing pipelines are often automated using frameworks like **Apache Spark** and **TensorFlow Data Services**, allowing scalable, real-time data preparation. By optimizing preprocessing and feature engineering, IDS models become not only more accurate and efficient but also more resilient to evolving cyber threats and adaptable to the dynamic nature of cloud environments.

Hybrid and Ensemble Learning Approaches:

Hybrid and ensemble learning approaches have emerged as powerful strategies for enhancing the performance and reliability of Artificial Intelligence (AI)-driven Intrusion Detection Systems (IDS) in cloud environments. Traditional machine learning models, when used individually, often face limitations in terms of accuracy, generalization, and robustness against diverse types of cyberattacks. To overcome these shortcomings, ensemble methods—such as bagging, boosting, and stacking—combine the predictions of multiple models to achieve superior detection results. Bagging (Bootstrap Aggregating), used in algorithms like Random Forest, reduces variance by training multiple models on different subsets of data and averaging their outputs. Boosting, as employed in methods like Gradient Boosting and AdaBoost, focuses on sequentially training weak learners by giving more weight to misclassified instances, thus improving the overall predictive power. Stacking, a more advanced technique, combines several base learners and uses a meta-learner to optimize final predictions, creating a layered architecture that enhances accuracy and reduces false alarms.

In addition to ensemble learning, hybrid intrusion detection systems integrate multiple detection methodologies—typically signature-based, anomaly-based, and specification-based models—to achieve both precision and adaptability. Signature-based components excel at identifying known threats using predefined attack patterns, while anomaly-based models detect deviations from normal behavior to capture zero-day or unknown attacks. Specification-based detection bridges the gap by defining legitimate behavior rules for applications and flagging activities that violate them. By combining these approaches, hybrid IDS ensures broader coverage of threat landscapes and minimizes blind spots that standalone models might miss.

The integration of hybrid and ensemble learning in cloud-based IDS is particularly beneficial in handling the heterogeneity and scale of cloud traffic, where no single model can efficiently address all types of attacks. For instance, ensemble systems can use supervised models for known threats and unsupervised clustering algorithms for unknown anomalies simultaneously. Furthermore, modern hybrid frameworks employ deep ensemble learning, where deep neural networks are combined with decision trees or probabilistic models to capture both low-level and high-level data representations. These approaches not only enhance detection accuracy but also improve resilience against adversarial attacks, as multiple models can cross-verify suspicious behaviors. Overall, hybrid and ensemble learning architectures represent the next stage in the evolution of IDS, providing a balanced blend of accuracy, adaptability, and interpretability essential for securing complex and dynamic cloud ecosystems.

Real-Time Intrusion Detection and Response Systems:

Real-time Intrusion Detection and Response Systems (IDS) have become indispensable in modern cloud computing environments, where cyber threats evolve rapidly, and even a few seconds of delay in detection can result in massive data breaches or service disruptions. Traditional batch-based IDS approaches are no longer sufficient, as they analyze static datasets after an event has occurred. In contrast, real-time IDS leverage streaming data analytics, online learning models, and automated response mechanisms to detect and neutralize threats as they unfold. These systems continuously process live network traffic, event logs, and system metrics, enabling instant anomaly identification. Technologies such as Apache Kafka, Spark Streaming, and AWS Kinesis are often used to handle massive data streams efficiently, ensuring low-latency analysis in distributed cloud infrastructures.

AI and machine learning play a crucial role in these systems by enabling online learning, where models update their parameters dynamically as new data arrives. This continuous learning capability ensures that IDS remain effective even against emerging attack patterns or behavioral shifts in network usage. For example, a real-time IDS may detect abnormal API calls, unauthorized data exfiltration, or command injection attempts within milliseconds, triggering an automated security response. The integration of Reinforcement Learning (RL) allows these systems to optimize their decision-making processes—choosing the best possible action (e.g., isolating a container, blocking an IP address, or throttling suspicious traffic) based on previous outcomes and current threat context.

The fusion of IDS with Security Information and Event Management (SIEM) platforms—such as Splunk, IBM QRadar, or ArcSight—enables comprehensive threat visibility across multiple layers of the cloud infrastructure. SIEM tools aggregate and correlate data from diverse sources, while AI-driven IDS enhance these platforms with predictive threat intelligence and automated remediation. The integration of automated incident response systems ensures that detected threats are mitigated immediately without human intervention, significantly reducing the mean time to detect (MTTD) and mean time to respond (MTTR).

real-time IDS act as the nervous system of cloud security, continuously sensing, analyzing, and responding to anomalies. The incorporation of AI, big data analytics, and automation has transformed them from passive monitoring tools into active, intelligent defense agents capable of adapting to dynamic environments. As cloud ecosystems continue to scale, real-time IDS will remain central to achieving zero-trust architectures and ensuring continuous protection against sophisticated cyber threats.

Energy-Efficient and Resource-Aware IDS Models:

Energy-efficient and resource-aware Intrusion Detection Systems (IDS) are an emerging necessity in large-scale cloud environments, where computational cost, power consumption, and latency directly affect operational efficiency. Traditional AI-driven IDS models—particularly those based on deep learning—require substantial computing power for training and inference, leading to high energy demands and scalability issues. To address this, researchers and engineers are focusing on optimization techniques that minimize resource utilization without compromising detection accuracy. Model compression methods such as quantization, pruning, and knowledge distillation are increasingly employed to reduce model size and computational complexity. Quantization converts high-precision parameters (like 32-bit floating points) into lower-bit representations (e.g., 8-bit), while pruning eliminates redundant neurons and connections from neural networks, resulting in lighter, faster models that maintain performance integrity.

Another promising approach is edge-AI deployment, which offloads portions of the intrusion detection workload to edge nodes or gateways located closer to data sources. This distributed design significantly reduces the data transmission overhead to central servers and minimizes latency in detecting attacks. Edge-based IDS can perform initial anomaly filtering, sending only critical alerts or aggregated features to centralized AI models for further analysis. This hybrid edge-cloud model improves both energy efficiency and scalability, especially in Internet of Things (IoT)-integrated cloud networks, where devices continuously generate massive data streams. Additionally, energy-efficient inference frameworks—such as TensorFlow Lite, PyTorch Mobile, and NVIDIA TensorRT—enable real-time IDS deployment on low-power hardware while maintaining high accuracy.

Cloud providers are also integrating dynamic resource allocation algorithms that adjust computing power based on real-time workload, ensuring that IDS models only consume resources proportional to network traffic intensity. Techniques like asynchronous processing, batch normalization, and caching of feature representations further optimize performance by reusing computations and reducing redundancy. Moreover, green AI principles are being adopted to measure and minimize the carbon footprint associated with model training and operation.

Ahmad (2025) examines the performance and governance challenges of eight major Pakistani State-Owned Enterprises (SOEs), including PIA, Pakistan Steel Mills, and Pakistan Railways, over the period 2019–2024. Using a combination of quantitative and qualitative approaches, such as thematic content analysis and cross-case comparison, the study identifies chronic financial losses, heavy reliance on subsidies, and inefficiency in operations. Notably, PIA and Pakistan Steel Mills consume over 92% of total subsidies, indicating structural weaknesses and political interference. Ahmad highlights that reforms like privatization, public-private partnerships, and professionalized governance are critical to restoring public trust, enhancing transparency, and achieving sustainable and accountable public sector management in Pakistan.

Ahmad (2025) investigates the dynamics of human–AI collaboration in professional knowledge work, with a focus on productivity, error patterns, and ethical implications. Participants were assigned to human-only, AI-assisted, and optional AI-only task groups performing activities such as writing, summarization, decision-support, and problem-solving. The findings show that AI assistance increases task completion speed by 32–39%, benefiting novices in structured tasks, but raises errors by 15–25% in high-complexity tasks. Ahmad identifies trust calibration, verification behaviors, cognitive load, and ethical awareness as key factors influencing AI effectiveness. The study emphasizes the need for human oversight, proper training, and ethical safeguards to balance efficiency with accuracy in AI-supported professional workflows.

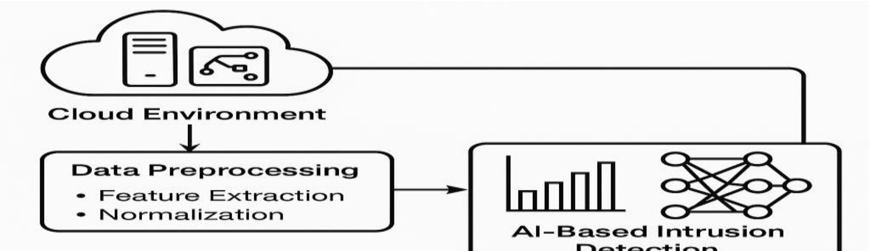

Summary

AI-driven intrusion detection systems represent a paradigm shift in cloud security. By harnessing machine learning and deep learning, these systems achieve superior detection accuracy, adaptability, and real-time response capabilities compared to traditional IDS. They empower cloud infrastructures with the intelligence to detect unknown threats and respond autonomously. However, the deployment of these systems must be carefully managed to address computational demands, data privacy, and explainability. Future research should focus on integrating federated learning, blockchain, and explainable AI to create transparent, scalable, and trustworthy intrusion detection systems for the cloud era.

References

Ahmed, M., Mahmood, A. N., & Hu, J. (2016). A survey of network anomaly detection techniques. Journal of Network and Computer Applications, 60, 19–31.

Buczak, A. L., & Guven, E. (2016). A survey of data mining and machine learning methods for cyber security intrusion detection. IEEE Communications Surveys & Tutorials, 18(2), 1153–1176.

Javaid, A., Niyaz, Q., Sun, W., & Alam, M. (2016). A deep learning approach for network intrusion detection. IEEE Trustcom/BigDataSE/ISPA, 190–197.

Alom, M. Z., et al. (2019). Intrusion detection using deep belief networks. Journal of Information Security and Applications, 48, 102–110.

Moustafa, N., & Slay, J. (2015). UNSW-NB15 dataset for network intrusion detection systems. Military Communications and Information Systems Conference, 1–6.

Shone, N., Ngoc, T. N., Phai, V. D., & Shi, Q. (2018). A deep learning approach to network intrusion detection. IEEE Transactions on Emerging Topics in Computational Intelligence, 2(1), 41–50.

Tang, T. A., et al. (2018). Deep learning approaches for network intrusion detection: A review. IEEE Access, 6, 38371–38394.

Li, J., Chen, J., & Zhao, X. (2020). AI-based anomaly detection in cloud computing. Future Generation Computer Systems, 108, 128–142.

Zhang, Y., & Zhou, Y. (2021). Federated learning for cloud intrusion detection. IEEE Transactions on Cloud Computing, 9(3), 1050–1063.

Kumar, R., & Kaur, P. (2022). Hybrid machine learning models for intrusion detection. Computer Networks, 212, 108–124.

Mehmood, S., et al. (2023). Privacy-preserving AI models for cloud security. Security and Privacy, 6(2), e173.

Tariq, H., et al. (2024). Blockchain-integrated intrusion detection frameworks. Journal of Cybersecurity and Digital Trust, 7(4), 223–240.

Ahmad, N. R. (2025). *Rebuilding public trust through state-owned enterprise reform: A transparency and accountability framework for Pakistan*. Punjab Sahulat Bazaars Authority (PSBA), Lahore, Pakistan. https://doi.org/10.24088/IJBEA-2025-103004

Ahmad, N. R. (2025). *Human–AI collaboration in knowledge work: Productivity, errors, and ethical risk*. https://doi.org/10.52152/6q2p9250