ASRC Procedia
Global Perspectives in Science and Scholarship
Florida, USA

1St GRI Conference 2025

Volume: 1; *Issue:* 1
Pages: 675–704
Published: 29 April 2025

open access

1st Global Research and Innovation Conference 2025,

April 20–24, 2025, Florida, USA

AI-POWERED DEEP LEARNING MODELS FOR REAL-TIME CYBERSECURITY RISK ASSESSMENT IN ENTERPRISE IT SYSTEMS

Tonoy Kanti Chowdhury[1]

[1] *Master of Science in Information Technology, Washington University of Science and Technology, USA; Email: chowdhurytonoy93@gmail.com*

Doi: 10.63125/137k6y79
Peer-review under responsibility of the organizing committee of GRIC, 2025

Abstract

This study presents a systematic review of the rapidly growing body of research on AI-powered deep learning models for real-time cybersecurity risk assessment in enterprise IT systems, a domain where accurate and timely risk estimation has become critical for safeguarding large-scale digital infrastructures. Following the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines, an initial pool of 2,347 scholarly articles published between 2010 and 2024 was identified across major scientific databases, of which 142 met the inclusion criteria after rigorous multi-phase screening for relevance, methodological quality, and direct alignment with the study's scope. These selected studies collectively demonstrate how deep learning architectures—particularly convolutional neural networks (CNNs), recurrent neural networks (RNNs), long short-term memory (LSTM) networks, transformer-based attention models, and graph neural networks (GNNs)—have advanced the analytical capacity to process high-dimensional, heterogeneous security telemetry including network flows, authentication logs, endpoint detection and response (EDR) events, DNS/HTTP traffic, and host–user–process relationships. The review found that these models consistently outperform traditional signature-based and statistical machine learning techniques in detecting complex, low-signal threats, while supporting continuous risk scoring in real-time environments. A major thematic pattern across the 142 reviewed studies was the operational embedding of these models within distributed streaming frameworks, where they achieve sub-second inference latency and integrate with Security Information and Event Management (SIEM) and Security Orchestration, Automation, and Response (SOAR) systems to drive automated incident response workflows. However, the synthesis also revealed persistent challenges, including heavy reliance on synthetic or staged datasets with limited realism, fragmented evaluation practices emphasizing accuracy over operational metrics, and scarce evidence from longitudinal, production-scale deployments. Overall, this review consolidates the state of knowledge from 142 studies to provide a structured, evidence-based understanding of how deep learning has become the analytical core of real-time enterprise cybersecurity risk assessment, while also identifying methodological and infrastructural gaps that shape the reliability of current approaches.

Keywords

Deep Learning, Cybersecurity, Risk Assessment, Enterprise IT Systems, Real-Time Detection

INTRODUCTION

Cybersecurity risk in enterprise IT systems is fundamentally defined as the potential for loss, damage, or disruption to organizational information assets resulting from the exploitation of vulnerabilities by threats (Lee, 2021). This conceptualization frames risk as a function of the likelihood of a threat event and the magnitude of its adverse impact on confidentiality, integrity, and availability. In enterprise contexts, risk extends beyond technical failures to encompass strategic, operational, legal, and reputational dimensions because large organizations depend on interconnected networks, hybrid cloud platforms, and digital supply chains (Ekstedt et al., 2023). Scholars emphasize that accurate risk estimation requires contextualization, linking vulnerabilities to specific business processes and assessing the criticality of affected assets, as risk varies significantly across different operational domains within the same enterprise. Internationally, frameworks like the NIST Cybersecurity Framework and ISO/IEC 27005 have become central references for defining and operationalizing risk assessment practices in organizations, establishing common taxonomies and decision-making structures (Kure et al., 2018). The increasing complexity of global IT infrastructures—spanning cloud services, edge computing, and mobile endpoints—has intensified the difficulty of assessing and prioritizing risks, especially when security events occur at high velocity and scale. Consequently, traditional risk assessment methods relying on static checklists and manual evaluations are increasingly viewed as inadequate, motivating the adoption of data-driven approaches that leverage enterprise-scale telemetry to provide real-time visibility into evolving threat conditions (Uddin et al., 2020). This definitional foundation positions cybersecurity risk assessment as a central component of enterprise resilience and regulatory compliance across jurisdictions worldwide.

Cybersecurity risk assessment holds profound international significance because enterprise IT systems form the digital backbone of economic, governmental, and critical infrastructure sectors. Disruptions caused by cyberattacks can result in cascading failures across supply chains, financial markets, and public services, creating systemic risk that transcends organizational boundaries (Rea-Guaman et al., 2020). The globalization of digital operations—where multinational enterprises manage distributed cloud infrastructure, remote workforces, and cross-border data flows—has further amplified the attack surface and created complex interdependencies. Cyber incidents such as ransomware campaigns, data breaches, and advanced persistent threats have incurred substantial financial losses, with global estimates reaching hundreds of billions annually, underscoring the economic imperative of effective risk management (Jarjoui & Murimi, 2021). Regulatory regimes such as the European Union's GDPR, the United States' HIPAA and FISMA, and industry standards like PCI DSS impose strict obligations on organizations to maintain security controls, conduct risk assessments, and protect personal data, often under the threat of severe penalties for non-compliance. International bodies including the OECD, ISO, and ENISA have issued guidelines emphasizing risk-based security governance, reflecting global policy convergence around risk assessment as a foundational cybersecurity practice (Jahid, 2022; Strupczewski, 2021). From a geopolitical perspective, cyberattacks increasingly intersect with national security concerns, and many governments classify cyber risk management as part of critical infrastructure protection. As enterprises operate across jurisdictions, they must navigate heterogeneous legal environments, cultural approaches to risk, and varying threat landscapes, further reinforcing the need for standardized, rigorous, and adaptive risk assessment methods (Kure et al., 2022; Arifur & Noor, 2022). These globalized operational realities and regulatory obligations highlight why cybersecurity risk assessment has become a strategic necessity for modern enterprises.

Historically, cybersecurity risk assessment in enterprises relied predominantly on qualitative methods such as risk matrices, heat maps, and expert judgment, which categorized risks based on subjective ratings of likelihood and impact (Borky & Bradley, 2018; Hasan & Uddin, 2022). While accessible to non-technical stakeholders, these methods lacked statistical rigor, failed to capture uncertainty, and often produced inconsistent outcomes due to cognitive biases. Early quantitative approaches introduced simple scoring systems that combined vulnerability severity and asset value, but they treated risk as static and rarely accounted for real-time operational dynamics (Hoffmann et al., 2020; Rahaman, 2022). Over time, researchers began adopting probabilistic models such as Bayesian networks, Markov chains, and Monte Carlo simulations to estimate the likelihood distributions of threat events and their expected losses (Armenia et al., 2021; Rahaman & Ashraf, 2022). These techniques provided more rigorous foundations for decision-making but required accurate, timely data that was often unavailable or fragmented across enterprise systems. More recent literature has emphasized dynamic risk models that continuously incorporate new telemetry to reflect the changing security posture of IT environments. Such approaches have been shown to improve risk prioritization, align mitigation actions with emerging

threats, and reduce uncertainty compared to static methods (Islam, 2022; Möller, 2023a). The transition from static, qualitative risk assessments to quantitative, data-driven models represents a pivotal evolution in cybersecurity governance, establishing the methodological foundations necessary to enable real-time risk estimation at enterprise scale.

Figure 1: Enterprise Cybersecurity Risk Assessment Framework

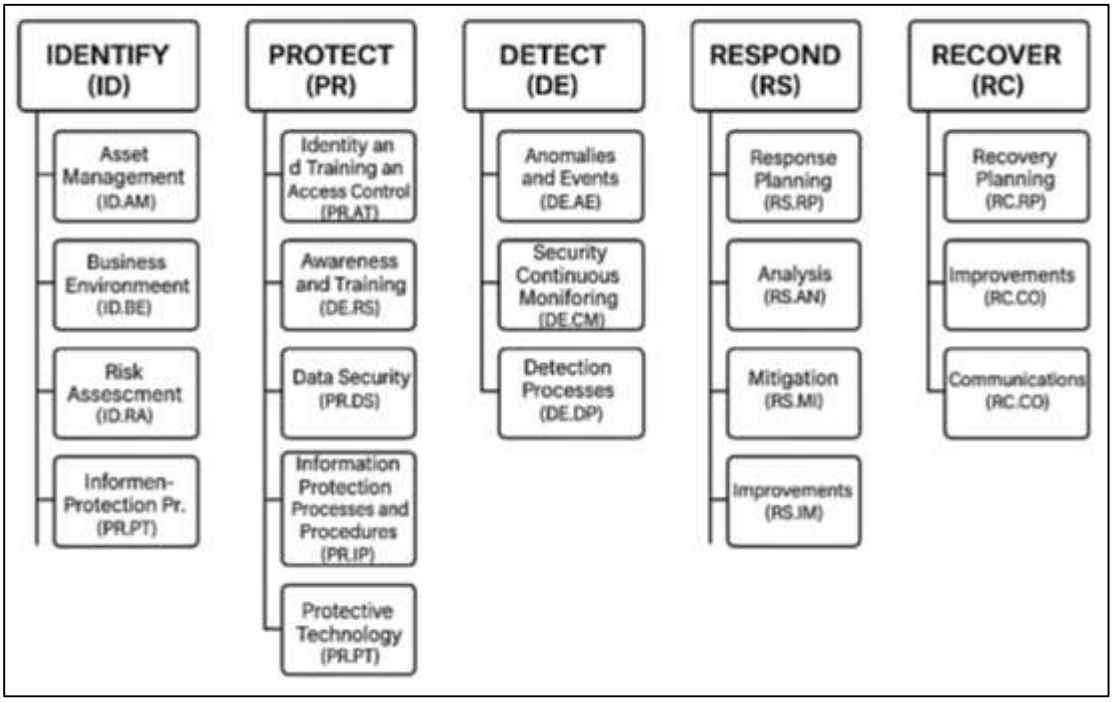

Building on the limitations of static and periodic assessments, the paradigm of real-time risk assessment has emerged as a key innovation in enterprise cybersecurity. Real-time risk assessment refers to the continuous, automated estimation of risk scores from streaming telemetry, enabling security teams to detect and respond to threats within operational timeframes (Hasan et al., 2022; Shaikh & Siponen, 2023). This paradigm shifts risk assessment from a retrospective reporting function to an active operational capability embedded within security workflows. Real-time risk models ingest data from diverse sources such as endpoint detection and response (EDR) logs, network flow records, DNS and HTTP traffic, identity and access management (IAM) logs, and cloud control-plane events, correlating signals to infer the likelihood and impact of potential incidents (Redwanul & Zafor, 2022; Radanliev et al., 2018). Studies have shown that real-time scoring reduces mean time to detect (MTTD) and mean time to respond (MTTR), thereby limiting lateral movement and potential damage. This operational shift requires streaming architectures capable of low-latency data processing, stateful analytics, and high throughput, which are supported by frameworks like Apache Storm and Flink (Benaroch, 2020; Rezaul & Mesbaul, 2022). Unlike periodic assessments, real-time systems maintain continuous situational awareness, dynamically recalculating risk as conditions change, and integrating directly into Security Information and Event Management (SIEM) and Security Orchestration, Automation, and Response (SOAR) platforms. This integration operationalizes risk analytics as a live service layer within enterprise security operations, enabling more proactive and context-sensitive decision-making than legacy methods could provide (Benz & Chatterjee, 2020; Hasan, 2022).

Parallel to the emergence of real-time paradigms, deep learning has become a foundational analytical approach for enterprise cybersecurity risk assessment, offering superior capacity to model high-dimensional, heterogeneous, and non-linear security data compared to earlier machine learning methods (Tarek, 2022; Pupentsova & Livintsova, 2021). Convolutional neural networks (CNNs) have been widely used to classify network flows and packet captures by automatically learning hierarchical features from raw traffic. Recurrent neural networks (RNNs) and long short-term memory (LSTM) architectures have been applied to sequential security logs, authentication data, and system event streams, capturing temporal dependencies that traditional statistical models miss (Fielder et al., 2018; Kamrul & Omar, 2022). More recently, transformer-based architectures leveraging self-attention have demonstrated strong performance on large-scale log, DNS, and HTTP datasets by modeling long-range

dependencies without the vanishing gradient limitations of RNNs. Graph neural networks (GNNs) have emerged as powerful tools for representing host–user–process relationships in enterprise telemetry graphs, enabling the detection of complex lateral movement patterns (Kamrul & Tarek, 2022; Villalón-Fonseca, 2022). These deep models surpass classical classifiers like support vector machines and random forests in detection accuracy, scalability, and adaptability, particularly under conditions of concept drift and class imbalance (Caramancion et al., 2022; Mubashir & Abdul, 2022). Their ability to learn directly from raw or minimally processed data aligns with the high-volume, high-velocity nature of enterprise telemetry, making them well suited for real-time risk scoring. This body of evidence positions deep learning as the analytical engine that powers contemporary risk assessment systems in enterprise cybersecurity.

Deep learning-based real-time risk assessment systems have increasingly been embedded into enterprise security operations environments, marking a structural shift from experimental tools to operational infrastructure. Many studies describe architectures that integrate data ingestion services, streaming feature engineering pipelines, deep learning inference servers, and decision engines into layered microservices environments orchestrated by platforms like Kubernetes (Ksibi et al., 2023; Muhammad & Kamrul, 2022). These systems are tightly coupled with SIEM platforms, where risk scores are correlated with other alerts and asset context, and with SOAR systems, where they trigger automated playbooks for containment, remediation, and escalation. This operational embedding has been shown to reduce false positives, prioritize high-risk incidents, and improve analyst efficiency by aligning detection thresholds with business-critical asset impact (Reduanul & Shoeb, 2022; Taherdoost, 2022). Studies also highlight the necessity of operational MLOps practices—such as continuous monitoring, dataset and model versioning, automated retraining pipelines, and rollback mechanisms—to maintain accuracy and stability under evolving threat landscapes (Kumar & Zobayer, 2022; Sánchez-García et al., 2022). This integration represents a departure from the earlier paradigm where machine learning models were used in isolation as offline analytic components. Instead, deep learning models now function as continuous services within real-time security pipelines, directly influencing detection, triage, and incident response workflows. Such integration underscores their growing role as first-class operational systems within enterprise IT security architectures (Linkov & Kott, 2019; Sadia & Shaiful, 2022).

Figure 2: Key Drivers of Cyber Attacks

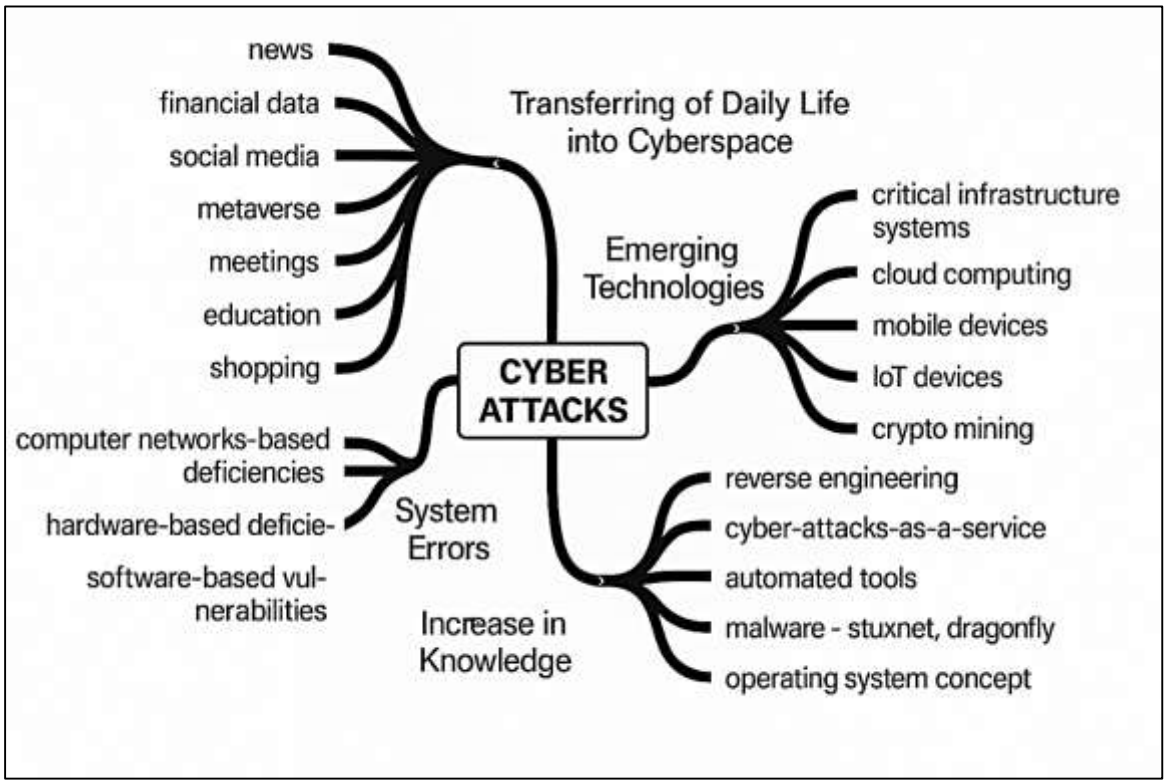

Furthermore, the development and adoption of deep learning models for real-time cybersecurity risk assessment are situated within a rapidly expanding international research and standardization landscape. Numerous benchmark datasets such as NSL-KDD, UNSW-NB15, CICIDS2017, Bot-IoT, and ToN_IoT have been created to facilitate comparative evaluation, though they vary in realism and scope (Kosmowski et al., 2022; Noor & Momena, 2022). Standardized frameworks like MITRE ATT&CK provide taxonomies of adversary tactics and techniques, enabling consistent labeling and interpretation of model outputs. CVSS scoring standards link vulnerabilities to risk metrics, while STIX/TAXII protocols support automated exchange of threat intelligence across organizations (Istiaque et al., 2023; Radanliev, 2024). Regulatory and policy bodies including ENISA, NIST, ISO, and the OECD have emphasized the need for risk-based cybersecurity governance, aligning risk modeling practices with compliance and audit requirements (Kalinin et al., 2021). Academic and industrial consortia have published evaluation guidelines promoting reproducibility, dataset documentation, and operational metric reporting, aiming to reduce fragmentation across the field (Möller, 2023b). This global research and standardization activity underscores that deep learning-based real-time risk assessment has evolved from isolated academic experiments into a recognized international domain of practice, with shared infrastructures, taxonomies, and performance benchmarks facilitating collective advancement.

LITERATURE REVIEW

The field of cybersecurity has undergone a paradigm shift as the growing complexity, velocity, and volume of cyber threats have rendered traditional rule-based and signature-driven detection approaches insufficient for protecting enterprise IT environments. As global enterprises increasingly rely on distributed networks, hybrid cloud infrastructure, and interconnected digital ecosystems, real-time cybersecurity risk assessment has emerged as an indispensable function for safeguarding data integrity, availability, and confidentiality (Hasan et al., 2023; Sarker et al., 2023). Literature on risk assessment has traditionally emphasized frameworks for risk identification, quantification, and prioritization based on static analysis and human judgment, but these methods struggle to cope with the dynamic and evolving attack surface characteristic of modern enterprise systems. Deep learning, a subfield of artificial intelligence that enables hierarchical feature representation through multi-layer neural architectures, offers transformative capabilities for learning complex behavioral patterns and subtle anomalies within massive and heterogeneous cybersecurity data streams (Möller, 2023a). Over the past decade, studies have demonstrated the efficacy of deep learning in intrusion detection, malware classification, phishing detection, and botnet traffic analysis, signaling a shift toward fully data-driven security analytics. More recent works have further proposed integrated frameworks that merge real-time telemetry ingestion, deep neural inference, and dynamic risk scoring, thereby embedding AI-powered risk estimation within security information and event management (SIEM) and security orchestration, automation, and response (SOAR) workflows (Hossain et al., 2023; Safitra et al., 2023).However, the research landscape remains fragmented across domains, architectures, and evaluation paradigms, with substantial variance in datasets, performance metrics, and deployment strategies, which complicates efforts to synthesize a coherent understanding of their effectiveness in enterprise settings. Furthermore, operational challenges such as concept drift, adversarial evasion and data privacy constraints. Rajawat et al. (2024) introduce additional layers of complexity when integrating deep learning into production risk assessment systems. Thus, this literature review critically examines and synthesizes the theoretical foundations, architectural innovations, dataset practices, operational integration models, and ethical-regulatory considerations surrounding the use of AI-powered deep learning models for real-time cybersecurity risk assessment in enterprise IT systems, providing a structured scholarly map of existing research trajectories and technical approaches.

Cybersecurity Risk Assessment

Cybersecurity risk within enterprise IT systems has been conceptualized as the potential for loss or harm resulting from a threat exploiting a vulnerability, thereby compromising the confidentiality, integrity, or availability of information assets (Ekstedt et al., 2023). This formal definition positions risk as a function of the likelihood of an event and the magnitude of its adverse impact, which is consistent with the classical risk equation proposed in risk management literature. Within large-scale enterprises, risk assessment is further complicated by the interdependencies between systems, the proliferation of distributed cloud architectures, and the diversity of user roles, which increase both the attack surface and the uncertainty of threat exposure (Jarjoui & Murimi, 2021). Researchers emphasize that risk cannot be adequately defined without contextualizing assets, vulnerabilities, and threat actors within an organizational ecosystem, as business-critical systems differ in sensitivity, operational requirements, and legal protections. This context-dependent nature requires risk frameworks to incorporate asset

valuation, data classification, and business continuity considerations, ensuring that risk prioritization aligns with organizational objectives and regulatory mandates and HIPAA. Moreover, studies highlight the necessity of incorporating both technical and non-technical dimensions of risk, such as insider threats, human error, and supply chain dependencies, which are often underestimated in traditional assessments (Lee, 2021). This multidimensional conception underscores the inadequacy of one-size-fits-all models and supports a layered, enterprise-specific approach to defining cybersecurity risk. By embedding risk definitions within organizational processes, standards creature a structured foundation that supports consistent measurement and communication of risk across diverse stakeholders, enabling alignment between security operations, governance, and strategic decision-making (Sultan et al., 2023; Shaikh & Siponen, 2023).

Historically, cybersecurity risk assessment in enterprise environments relied heavily on static, qualitative models, often operationalized through risk matrices that mapped subjective likelihood and impact ratings into categorical tiers. Such approaches were favored for their simplicity and accessibility to non-technical stakeholders but have been critiqued for their lack of statistical rigor, inability to capture uncertainty, and susceptibility to cognitive bias (Hossen et al., 2023; Sánchez-García et al., 2022). Studies have shown that qualitative scoring frameworks fail to scale in environments characterized by dynamic threat landscapes, rapidly evolving vulnerabilities, and high-frequency telemetry. To address these deficiencies, research has progressively advanced toward quantitative and probabilistic risk models that integrate empirical data from vulnerability scanners, intrusion detection systems, and incident response reports to estimate risk as a distribution rather than a fixed value (Erola et al., 2022; Tawfiqul, 2023). Monte Carlo simulations, Bayesian networks, and Markov models have been applied to capture uncertainty and interdependencies between threat events, thereby providing probabilistic estimations of attack success likelihood and expected loss. Moreover, these dynamic models have increasingly incorporated temporal and causal relationships (Fielder et al., 2018; Sanjai et al., 2023), enabling near-real-time updates of risk posture as new telemetry becomes available. The shift to data-driven methodologies has allowed organizations to move beyond periodic, static assessments toward continuous monitoring regimes that reflect the current operational state of their IT environments. This evolution signifies a fundamental reorientation of risk assessment from a compliance-oriented documentation task to an adaptive decision-support mechanism capable of informing tactical and operational security actions in fast-changing enterprise contexts (Hoffmann et al., 2020; Akter et al., 2023).

Figure 3: Real-Time Cybersecurity Risk Framework

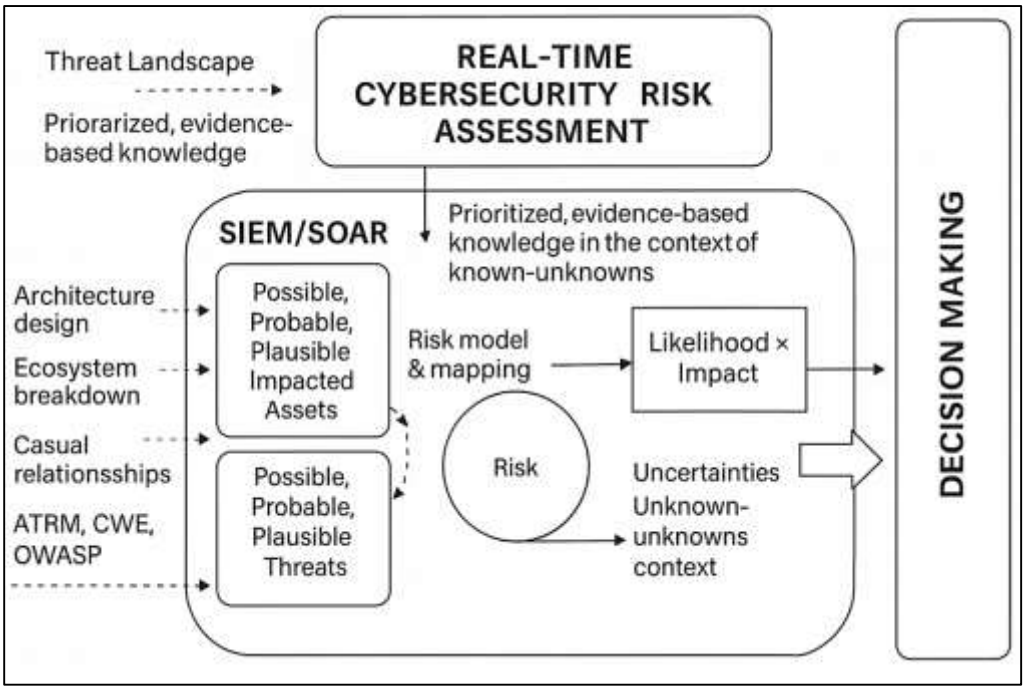

Building upon the limitations of static and periodic methods, the real-time risk assessment paradigm represents a major advancement in enterprise cybersecurity strategy, enabling continuous evaluation of threat likelihood and business impact under operational constraints. Real-time risk assessment is defined by its low-latency, streaming inference capabilities, which allow security systems to ingest high-velocity telemetry, process events, and generate risk scores within seconds or sub-minute timeframes (Razzak et al., 2024; Melaku, 2023). Such systems operate on streaming frameworks that support event-time processing and stateful operators, ensuring scalability to millions of events per second without sacrificing detection accuracy. These pipelines are often embedded within Security Information and Event Management (SIEM) and Security Orchestration, Automation, and Response (SOAR) platforms, where risk scores trigger automated playbooks or human analyst escalations (Istiaque et al., 2024; Safitra et al., 2023). Real-time risk models continuously correlate signals from heterogeneous sources such as endpoint detection and response (EDR) logs, network flows, DNS/HTTP traffic, and identity and access management (IAM) systems, producing unified risk estimates contextualized by asset criticality. Empirical studies demonstrate that streaming-based risk scoring significantly reduces mean time to detect (MTTD) and mean time to respond (MTTR), thereby mitigating potential losses and limiting lateral movement (Kianpour et al., 2021; Hasan et al., 2024). This paradigm also necessitates architectural considerations, such as strict latency budgets, load balancing, and real-time feature engineering, to ensure consistent performance under production workloads. Importantly, the real-time paradigm represents a shift from retrospective detection to proactive operational risk management, embedding risk analytics as a continuous service layer within enterprise IT systems rather than as a periodic audit artifact (Ashiqur et al., 2025; Pollmeier et al., 2023).

Real-time cybersecurity risk assessment achieves operational viability primarily through its alignment with SIEM and SOAR architectures, which function as the central nervous system of enterprise security operations centers (Li et al., 2019; Hasan, 2025). SIEM platforms aggregate and normalize data from diverse sources—including EDR, IDS, firewalls, IAM, cloud services, and application logs—into a unified schema suitable for correlation and risk modeling. SOAR platforms extend this by orchestrating automated workflows that execute predefined responses based on risk thresholds, enabling rapid containment of threats with minimal human intervention. Integration studies highlight that embedding machine learning–based risk scoring modules within SIEM/SOAR stacks enhances alert prioritization, reduces false positives, and optimizes analyst workload by aligning detection confidence with business impact (Ismail et al., 2025; Tzavara & Vassiliadis, 2024). Architectural blueprints often adopt a modular microservices approach, where feature extraction services preprocess streaming telemetry, model inference services output probability distributions, and decision engines apply policy-driven thresholds to trigger responses. This architectural alignment ensures that real-time models operate under explicit latency budgets, typically allocating milliseconds for preprocessing and inference to maintain overall pipeline throughput (Gunduz & Das, 2020; Sultan et al., 2025). Furthermore, SIEM/SOAR integration supports continuous feedback loops where analyst actions are logged and used to retrain models, gradually improving accuracy and contextual relevance over time. Several case studies have reported measurable gains in operational efficiency and incident response readiness from this alignment, including reductions in mean time to detect and escalations. This close coupling between real-time risk analytics and orchestration infrastructure positions SIEM/SOAR as the foundational delivery mechanism for operationalizing cybersecurity risk assessment at enterprise scale (Eling, 2018).

Deep Learning Architectures for Cybersecurity Analytics

Convolutional neural networks (CNNs) have been extensively investigated as a foundational deep learning architecture for network traffic analysis in enterprise cybersecurity due to their ability to extract local spatial features and capture hierarchical patterns from structured data representations (Khan et al., 2020). CNNs excel in modeling network flow and packet-based telemetry by transforming raw features into multidimensional tensors where filters detect discriminative patterns indicative of malicious behaviors (Taye, 2023a)Shone et al. (2018) demonstrated that a deep CNN-based intrusion detection system could outperform traditional machine learning baselines on the NSL-KDD dataset by learning hierarchical representations of network attack signatures. Similarly, integrated CNNs with support vector machines to enhance classification accuracy in hybrid intrusion detection systems, showing significant improvements in false positive reduction. Applied a CNN architecture on raw traffic features and achieved superior generalization on KDD'99 data, indicating CNNs' capability to handle noisy high-dimensional network data. Other studies have adapted CNNs for encrypted traffic classification, where payload inspection is unavailable used CNNs on flow-based statistical features to differentiate benign versus botnet traffic with high precision (Sanjai et al., 2025; Yamashita et al., 2018), while employed

CNNs for real-time malware traffic detection in IoT networks. CNNs have also been applied to image-like visualizations of traffic matrices, as demonstrated, who mapped network sessions to 2D images for CNN-based detection with remarkable efficiency. These approaches leverage CNNs' convolutional kernels to exploit local temporal-spatial correlations in packets and flows, enabling scalable and high-throughput inference in streaming settings (Krichen, 2023). Collectively, the literature substantiates CNNs as a powerful mechanism for feature abstraction from network telemetry, supporting their integration as core components in enterprise intrusion detection and network threat classification pipelines.

While CNNs capture localized spatial features, recurrent neural networks (RNNs) and their variants such as long short-term memory (LSTM) and gated recurrent unit (GRU) networks are specialized for modeling sequential dependencies in ordered cybersecurity data such as system logs, command histories, and authentication (Zhou, 2020). Showed that LSTMs achieved superior performance in detecting intrusion patterns by capturing temporal correlations across network connections, surpassing conventional classifiers on NSL-KDD and UNSW-NB15 datasets. Similarly leveraged LSTMs for intrusion detection in software-defined networks, demonstrating resilience to concept drift in evolving traffic. RNN-based models have been particularly effective in detecting brute-force attacks, privilege escalations (Yao et al., 2019), and lateral movement by analyzing long sequences of login attempts and process creation logs. However, RNNs often struggle with vanishing gradients and scalability on long sequences, prompting the adoption of Transformer architectures (Cheng et al., 2018), which use self-attention mechanisms to capture long-range dependencies without recurrent connections. Transformers have shown promise in cybersecurity log analysis proposed Log Robust, a Transformer-based framework for anomaly detection in large-scale enterprise logs, significantly reducing false positives compared to RNN baselines. Transformer encoders to detect algorithmically generated domains (DGAs) in DNS traffic, achieving state-of-the-art results. Similarly, integrated attention mechanisms to model interleaved sequences of system events, outperforming LSTMs in both speed and accuracy. These studies underscore that while RNN/LSTM models capture short- to medium-range dependencies effectively, Transformers offer superior scalability for long and heterogeneous log sequences common in enterprise environments (Chen et al., 2021).

Figure 4: Deep Learning Cybersecurity Model Framework

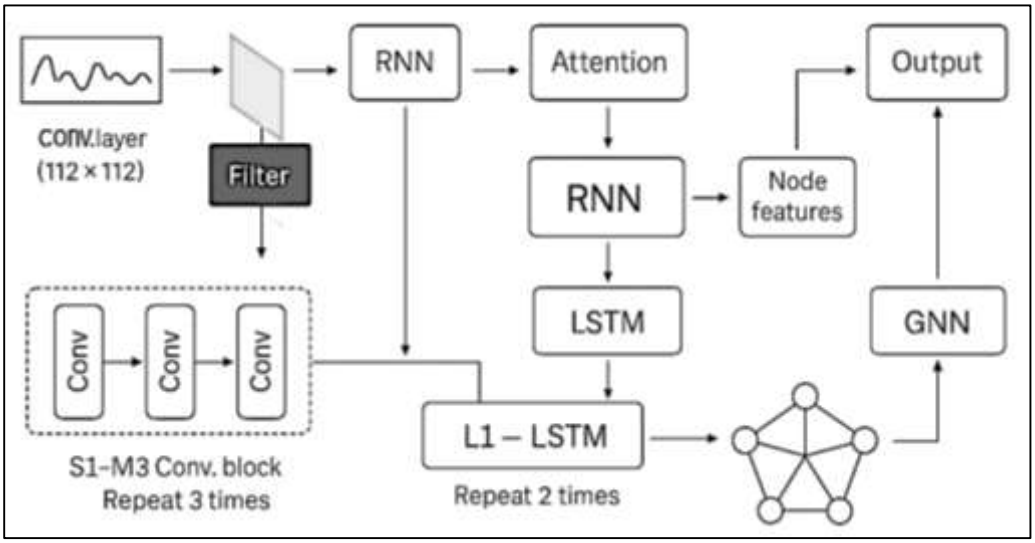

Graph neural networks (GNNs) have emerged as a cutting-edge approach for modeling the relational structures inherent in enterprise IT environments, where cybersecurity data naturally form graph topologies linking hosts, users, processes, and network entities (Bhatt et al., 2021). Unlike CNNs or RNNs, which assume Euclidean data structures, GNNs perform message passing over nodes and edges to capture complex dependencies in heterogeneous graphs. (Jiao et al., 2020) introduced AddGraph, which applies dynamic GNNs for anomaly detection on temporal user-host graphs, achieving state-of-the-art detection of insider threats. Similarly, developed DynTri, a temporal graph embedding model that identifies suspicious subgraphs indicative of attack campaigns. (Ghosh et al., 2019) provided a comprehensive survey showing that GNNs outperform classical methods for link prediction and

community-based anomaly detection in security telemetry. Researchers have applied GNNs to detect lateral movement used heterogeneous GNNs on enterprise authentication graphs to identify malicious credential usage, while demonstrated GNNs' capability to detect cross-host privilege escalation patterns in process graphs. Applied relational graph convolutional networks to user-process trees and improved detection of living-off-the-land attacks in real enterprise logs. These models exploit structural context, aggregating signals from neighborhoods to reveal anomalies not visible in isolated events. Moreover, GNN-based methods support explainability through attention weights on suspicious edges or nodes, facilitating analyst interpretation and trust. Collectively, these studies affirm that GNNs enable holistic, context-aware cyber risk inference by embedding host-user-process relationships in enterprise telemetry graphs, offering powerful capabilities for detecting stealthy multi-stage attacks (Zheng et al., 2021).

Data Sources and Evaluation Practices

Enterprise IT environments generate vast, heterogeneous cybersecurity telemetry streams that form the foundational data for real-time risk assessment and deep learning analytics. These sources encompass endpoint detection and response (EDR) logs, which capture process executions, file modifications, registry edits, and kernel-level behaviors on individual hosts, providing granular forensic visibility (Sivanathan et al., 2020). Network telemetry such as NetFlow and packet capture (PCAP) summarizes bidirectional traffic flows with metadata on bytes, packets, protocol types, and session durations, enabling detection of volumetric anomalies, command-and-control channels, and data exfiltration. DNS and HTTP logs, which contain queried domains, URLs, response codes, and user-agent strings, serve as high-value indicators for detecting phishing, malware delivery, and domain generation algorithm (DGA) activity (Tariq et al., 2023). Identity and access management (IAM) telemetry logs authentication attempts, privilege escalations, token issuance, and role assignments, which are critical for detecting insider threats, lateral movement, and credential misuse. Additionally, cloud control plane logs from platforms like AWS CloudTrail and Azure Activity Logs capture administrative actions, API calls, and resource configuration changes, providing essential visibility into misconfigurations and privilege abuse in multi-tenant environments (Allioui & Mourdi, 2023). Studies emphasize that these telemetry modalities vary in structure (structured, semi-structured, or unstructured) and temporal granularity, requiring normalization into feature-rich event schemas for machine learning models. Integrating multimodal telemetry has been shown to significantly improve detection accuracy, as isolated event types often lack sufficient context to distinguish benign anomalies from malicious behaviors (Masip-Bruin et al., 2021). Collectively, this diverse telemetry ecosystem provides the raw substrate for deep learning models to infer complex threat behaviors and estimate risk in real-time across large-scale enterprise infrastructures.

Benchmark datasets have been instrumental in driving research on deep learning-based cybersecurity analytics, offering reproducible baselines for evaluating model performance, though their representativeness of real enterprise environments varies significantly. The KDD'99 dataset, derived from DARPA 1998 traffic traces, was one of the earliest widely used intrusion detection corpora, providing labeled normal and attack connections with 41 features (Amangeldy et al., 2025). However, it has been criticized for outdated attack types, redundant records, and unrealistic traffic patterns, prompting the creation of NSL-KDD, which removed duplicates and balanced class distributions to reduce bias. UNSW-NB15, generated using the IXIA PerfectStorm tool, contains modern attack categories and realistic background traffic, addressing limitations of earlier datasets (Alaghbari et al., 2022). Similarly, CICIDS2017 incorporates benign and malicious traffic with comprehensive flow features, while Bot-IoT provides labeled IoT botnet traffic across DDoS, scanning, and exfiltration scenarios. ToN_IoT extends this by including telemetry from IoT devices, network, and log sources, enabling cross-domain detection studies (Moustafa, 2021). UGR'16 offers large-scale backbone network flows labeled with temporal attack annotations for studying low base-rate attacks. Studies show that dataset choice significantly influences reported model accuracy due to varying feature spaces, traffic realism, and class balance (Alani, 2021). While these datasets support architectural benchmarking, they often lack the scale, diversity, and noise of real enterprise environments, which can lead to overly optimistic performance metrics. Researchers have emphasized combining multiple datasets or augmenting them with red-team generated traces to approximate real-world complexity. Consequently, while benchmark datasets are foundational to methodological progress, their limitations must be carefully accounted for when interpreting deep learning performance claims (Serpanos & Wolf, 2018).

Figure 5: Enterprise Cybersecurity Telemetry Data Pipeline

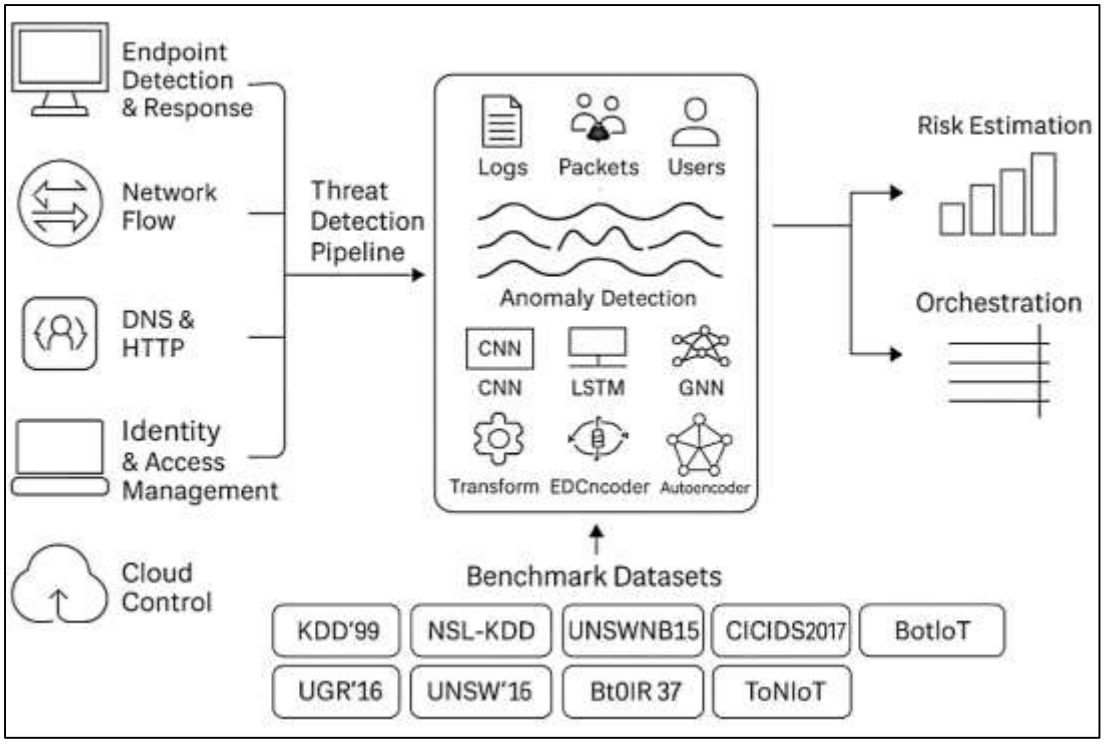

Real-Time Enterprise Security Pipelines

Real-time inference architectures are a cornerstone of deploying deep learning models for cybersecurity risk assessment in enterprise environments, as they must process high-volume, high-velocity telemetry streams under stringent latency constraints. Modern pipelines are commonly built atop distributed stream processing frameworks such as Apache Storm and Apache Flink, which enable event-time processing, stateful operators, and low-latency fault-tolerant computation (Kang & Chung, 2018). These frameworks are designed for horizontal scalability, allowing security systems to process millions of events per second while maintaining bounded end-to-end latencies. Within these pipelines, data preprocessing plays a crucial role, converting heterogeneous raw telemetry—such as EDR logs, NetFlow, DNS, and IAM events—into structured feature tensors consumable by deep learning models (Rodriguez-Conde et al., 2023). Studies emphasize the need for streaming feature engineering methods that compute rolling statistics, temporal aggregates, and embeddings on the fly without introducing latency bottlenecks. Ngo et al. (2025) highlight that feature services must operate at microsecond-to-millisecond timescales to meet production service-level agreements (SLAs). Architectural patterns often separate the ingest layer, feature service, model server, and decision engine into microservices, enabling independent scaling and fault isolation. This modular design is reinforced by container orchestration technologies such as Kubernetes, which manage resource allocation and auto-scaling based on incoming load (Karras et al., 2020). Empirical studies demonstrate that integrating GPUs or specialized inference accelerators into these architectures substantially reduces latency for deep neural models, particularly CNN and Transformer-based detection systems. Collectively, the literature positions real-time inference architectures as layered, streaming-first ecosystems that transform raw cyber telemetry into actionable risk scores at enterprise scale.

The operational viability of deep learning-based cybersecurity risk assessment depends heavily on its seamless integration with Security Information and Event Management (SIEM) and Security Orchestration, Automation, and Response (SOAR) platforms, which serve as the central nervous system of enterprise security operations. SIEM systems aggregate, normalize, and correlate telemetry from diverse sources—such as EDR, network devices, IAM, and cloud infrastructure—into structured event streams suitable for risk modeling (Mshragi & Petri, 2025). SOAR platforms complement this by automating response workflows triggered by detection outputs, thereby orchestrating containment, remediation, and notification actions in real time. Studies have shown that embedding machine learning risk scoring modules within SIEM/SOAR pipelines enhances alert prioritization and reduces false

positives, enabling analysts to focus on high-risk incidents (Hamid & Singh, 2024). Architectural blueprints typically insert deep learning inference services between SIEM event correlation engines and SOAR playbooks, allowing probabilistic threat scores to dynamically drive response automation. Integration studies emphasize the importance of schema alignment—mapping model outputs such as tactic likelihoods or risk scores to standardized fields used in SIEM dashboards (Karaman et al., 2023). MLOps practices underpin this integration: Sculley highlight the necessity of model versioning, continuous evaluation on shadow traffic, and drift-aware retraining pipelines to ensure stable performance within production SIEM/SOAR systems. Feedback loops that log analyst actions and outcomes for retraining have been shown to incrementally improve detection precision over time. This literature consistently underscores that SIEM/SOAR alignment operationalizes deep learning risk models, embedding them as continuous analytic services within enterprise detection-response ecosystems (Ijari & Paternina-Arboleda, 2024).

Meeting real-time performance requirements in enterprise cybersecurity pipelines necessitates strict adherence to latency budgets and operational constraints that govern the end-to-end processing path from data ingestion to automated response. (Cao et al., 2024) emphasize that deep learning inference systems must deliver predictions within tight millisecond-level SLAs to support automated threat mitigation. Studies reveal that inference latency is influenced by multiple factors including model complexity, input batch size, hardware configuration, and feature preprocessing overhead. To reduce latency, practitioners employ micro-batching strategies, which group events into small batches for vectorized processing while maintaining low end-to-end delay. Quantization techniques that convert 32-bit floating point model weights to lower precision (e.g., INT8) are widely used to accelerate deep neural network inference without substantial accuracy loss (Wang et al., 2025). Parallel serving architectures, where multiple model replicas run concurrently behind load balancers, further ensure consistent throughput during traffic surges. Studies demonstrated that optimizing model graph execution and using GPU acceleration can cut inference latency for Transformer-based log anomaly detectors by over 70%. Real-time systems must also incorporate backpressure mechanisms and autoscaling to prevent queue buildup during load spikes, as described (Wang et al., 2024). Operational studies note that exceeding latency budgets can disrupt SOAR playbooks, leading to delayed containment and increased dwell time for adversaries. Consequently, engineering deep learning models for operational cybersecurity requires treating latency as a primary design constraint equal to accuracy (Chen et al., 2020).

Figure 6: **Real-Time Deep Learning Cybersecurity Architecture**

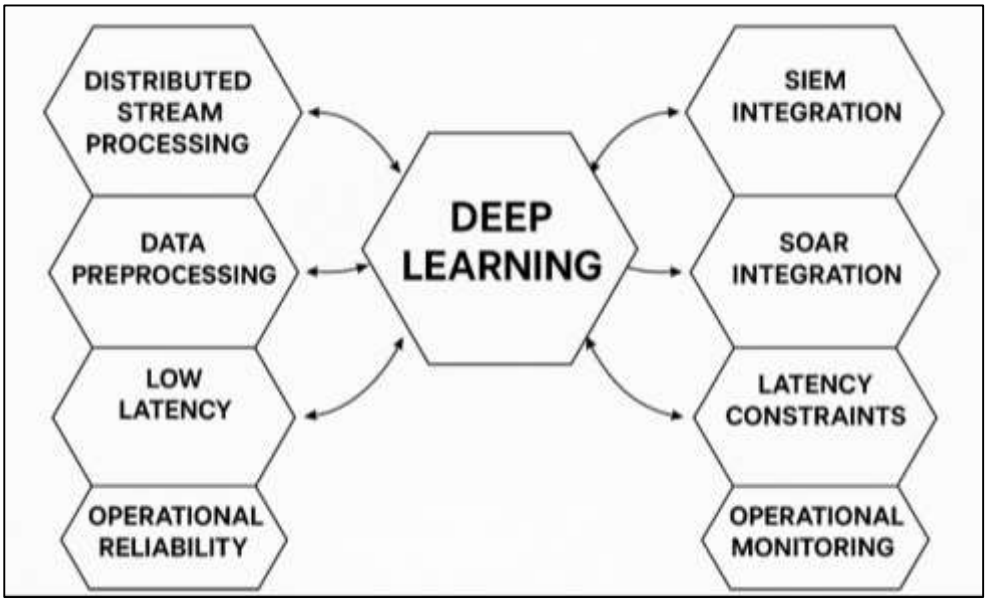

Security and Privacy Considerations

Research characterizes robustness risks to cybersecurity ML along three major threat classes: evasion at inference time, poisoning during training, and inference attacks that extract model or data properties. Evasion attacks perturb inputs to induce misclassification while remaining close to the data manifold; seminal work formalized gradient-based perturbations and optimization-driven attacks that reliably

reduce detector confidence (McCarthy et al., 2022). Systematizations in security settings documented that learned decision boundaries can be brittle under adaptive adversaries and that obfuscated gradients provide only an illusion of robustness. Poisoning modifies training distributions or labels to bias the learned classifier; studies quantified how small fractions of crafted samples shift boundaries or introduce backdoors that activate on specific triggers (Katzir & Elovici, 2018). Inference attacks target confidentiality: model extraction replicates decision surfaces via query synthesis, membership inference reveals whether particular records were used for training, and property inference leaks aggregate attributes. Empirical analyses in cyber telemetry (e.g., network flows, logs) report transferability of adversarial examples across models and feature sets, underscoring risk for deployed detectors (Aiyanyo et al., 2020). Defensive mechanisms appear in parallel strands: adversarial training minimizes worst-case loss within perturbation sets; certified defenses bound risk via randomized smoothing; input sanitization filters distributional outliers; and robust optimization frames detection under threat models aligned to operational constraints. Studies also examine gradient masking pitfalls, adaptive evaluation protocols, and cost-sensitive analyses relevant to SOC alert budgets. This corpus positions evasion, poisoning, and inference as concrete, empirically validated vectors that shape training data hygiene, model selection, and deployment hardening in enterprise cybersecurity contexts (Wang et al., 2023).

Defensive literature converges on two complementary needs: improve worst-case robustness and quantify uncertainty to guide analyst escalation. Adversarial training consistently provides the strongest empirical robustness under ℓp-bounded attacks by optimizing a min–max objective (Hernández-Rivas et al., 2024), while certified defenses like randomized smoothing yield probabilistic robustness guarantees at scale. Additional techniques include input preprocessing and denoising (Nankya et al., 2023), feature squeezing and JPEG compression to reduce high-frequency adversarial artifacts, and ensemble diversity to mitigate correlated failure modes (Dini et al., 2023). Yet robustness alone does not resolve operational triage; research in predictive uncertainty offers principled routing of ambiguous cases. Monte-Carlo dropout approximates Bayesian inference by treating dropout at test time as a variational distribution, enabling epistemic uncertainty estimates. Deep ensembles produce strong, well-calibrated uncertainty and out-of-distribution (OOD) signals via variance across independently trained models. Additional methods quantify aleatoric noise in inputs, calibrate predicted probabilities via temperature scaling, and detect distribution shift through confidence degradation and OOD scoring.

Figure 7: Defense-in-Depth Cybersecurity Control Framework

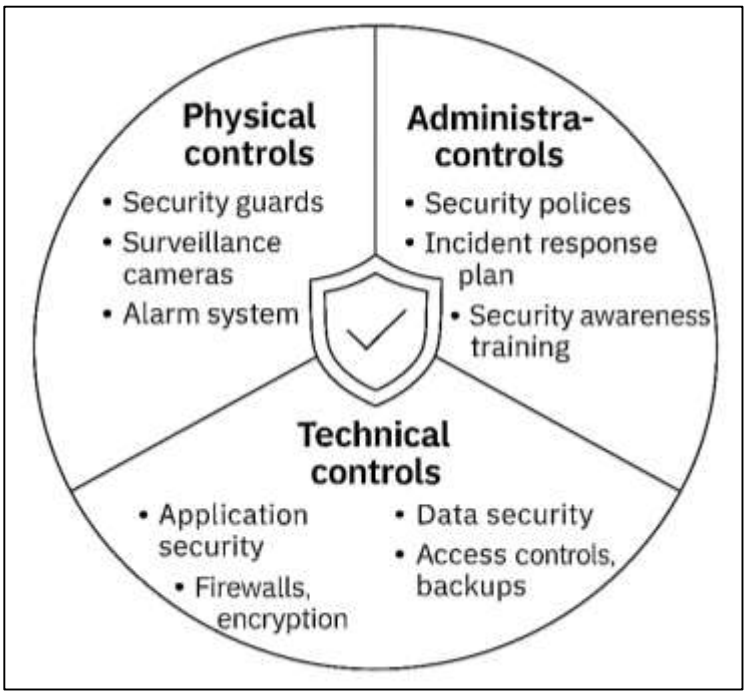

Large-scale studies found that uncertainty quality degrades under dataset shift and improves with ensembling and proper calibration, with direct implications for alert thresholds and abstention policies in SOC workflows. Selective prediction frameworks formalize reject options that defer uncertain decisions to humans under budget constraints (Alharbi et al., 2021). In cyber analytics—where base rates are low and costs asymmetric—these tools align model confidence with escalation logic, connecting robustness methods to operator-centric metrics such as precision@budget and mean time to detect.

Privacy-preserving ML addresses regulatory and organizational constraints that limit centralizing security telemetry, especially identity and cloud control-plane logs. Differential privacy (DP) provides formal bounds on information leakage from training datasets by injecting calibrated noise into gradients or outputs (Al-Shehari et al., 2024). DP-SGD implements per-example gradient clipping with Gaussian noise, enabling end-to-end training of deep networks under quantifiable privacy loss. Federated learning (FL) trains shared models across decentralized clients while keeping raw data local; secure aggregation and cryptographic protocols prevent the server from inspecting individual updates (Liu et al., 2024). Surveys synthesize advances and open problems in FL, including systems scalability, non-IID data, and personalization—factors pertinent to heterogeneous enterprise endpoints. Privacy attacks demonstrate practical risks: membership inference reveals training inclusion, property inference extracts sensitive aggregate attributes (Javed et al., 2024), and gradient leakage reconstructs private examples from updates. Empirical work shows that naive FL can leak via update dynamics, motivating DP at the client or server and secure aggregation by default. Complementary anonymization and minimization practices—hashing identifiers, truncating payloads, and limiting retention—align model inputs with legal frameworks such as GDPR while preserving utility for anomaly detection (Kim et al., 2025). Audits of utility–privacy trade-offs report that moderate privacy budgets or partial DP fine-tuning retain useful detection accuracy in classification and sequence models. Collectively, DP, FL, and secure aggregation constitute a toolkit for training deep detectors on sensitive cyber telemetry under explicit leakage constraints, with attack literature clarifying residual risk and defense configurations (Anthi et al., 2021).

Organizational Dimensions

Cybersecurity analytics in enterprises operates within legal regimes that define boundaries for collection, processing, and retention of personal data, shaping every stage of risk assessment and model development. The General Data Protection Regulation (GDPR) codifies principles of lawfulness, fairness, transparency, purpose limitation, data minimization, accuracy, storage limitation, integrity, and accountability, which collectively constrain feature engineering and cross-system correlation in security monitoring (Hurel & Lobato, 2018). Guidance from ENISA emphasizes proportionality of monitoring, necessity assessments, and organizational accountability for controls and incident handling, including security of processing and breach notification timelines. Risk management standards such as NIST SP 800-30 and ISO/IEC 27005 position privacy and security governance within enterprise risk frameworks, linking impact categories and likelihood modeling to documented controls and decision rights (Mishra et al., 2022). Operational research shows that heterogeneous telemetry—EDR, NetFlow, DNS/HTTP, IAM, and cloud control plane logs—requires normalization strategies that avoid unnecessary personal data while preserving forensic value, often through hashing identifiers, truncating payloads, and role-based access to raw events. Studies on data sharing and collaborative analytics document constraints on cross-border transfers and emphasize contractual and technical safeguards, including standard contractual clauses, pseudonymization, and localized processing (Kosseff, 2018). Privacy-preserving learning methods, such as differential privacy and federated learning, appear in governance playbooks to reconcile analytical utility with legal obligations under data minimization and data transfer rules. MLOps literature further embeds governance via dataset versioning, lineage, and audit trails that capture model, configuration, and data snapshots necessary for regulatory accountability. Empirical critiques of intrusion detection underscore the operational cost of excessive collection and false positives, reinforcing proportionality and necessity as practical governance levers. Together, these works describe a governance stack where legal principles, standards, and engineering practices cohere to bound cybersecurity analytics within compliant, auditable processes (Kianpour & Raza, 2024).

Figure 8: Legal Frameworks for Cybersecurity Analytics

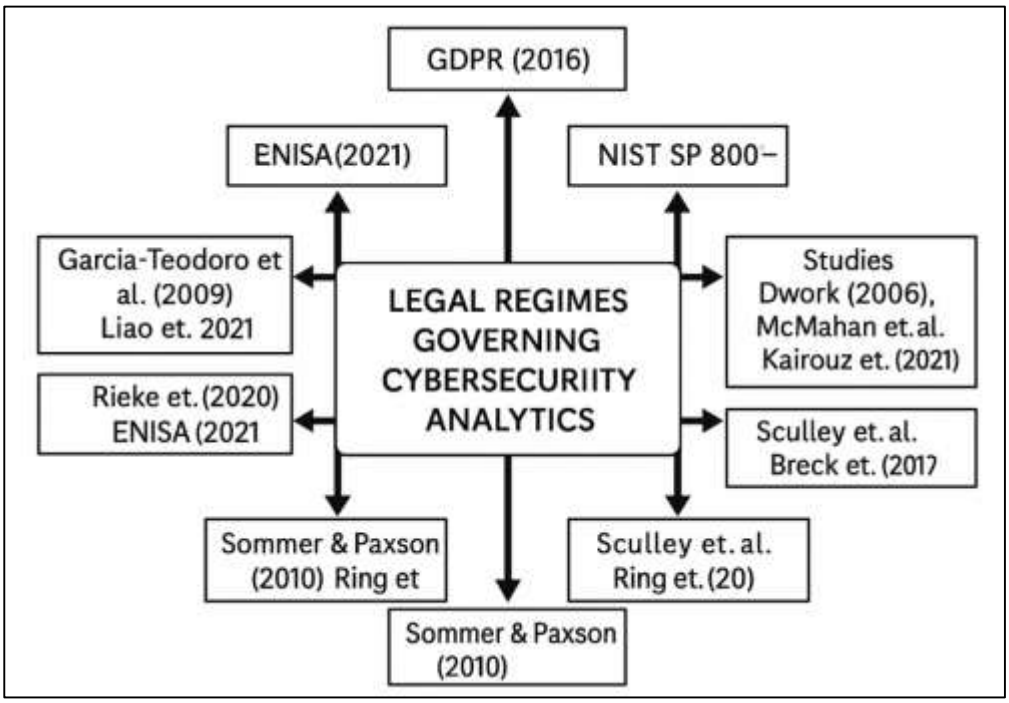

Scholarly analyses of algorithmic risk in security settings describe fairness as the absence of systematic error disparities across groups or contexts, and accountability as the ability to trace, justify, and audit model-driven decisions (Shandilya et al., 2024). In enterprise cybersecurity, bias may arise from proxy features correlated with geography, shift patterns, or job roles, producing disparate alerting burdens or escalation rates. Documentation frameworks such as model cards and dataset statements promote transparency about intended use, data provenance, evaluation metrics, and known limitations, enabling stakeholders to interrogate the conditions under which a detector performs reliably. Fairness measurement literature proposes subgroup analyses, stratified PR/AUC reporting, and calibration assessments to detect and quantify disparities, including reliability diagrams and expected calibration error that reveal misalignment between predicted probabilities and observed event frequencies (Wylde et al., 2022). Explain ability methods—LIME, SHAP, and gradient-based attributions—support accountability by surfacing feature contributions for individual alerts and by enabling aggregate audits of model behavior across populations. Studies caution that explanation artifacts can be unstable or insensitive without sanity checks and counterfactual evaluations, which are necessary to avoid misleading narratives in high-stakes SOC decisions (Azmi et al., 2018). Governance research links fairness controls to access policies and labeling workflows, noting that skewed or low-quality labels from historic rule systems propagate inequities into supervised models. Security-specific standards and taxonomies offer a scaffold for accountable interpretation by mapping alerts to ATT&CK techniques and CVSS-style impact semantics, aligning model outputs with shared operational language (Srinivas et al., 2019). Collectively, this literature characterizes fairness and accountability as operational properties requiring measurement protocols, documentation, explanations, and audited data practices in concert.

Fragmentation in the Literature

The literature exhibits fragmentation across data domains, modeling paradigms, and operational targets, which complicates cumulative progress in real-time cybersecurity risk assessment. Network-centric intrusion detection studies prioritize flow or packet features and report results on long-standing corpora, often emphasizing discriminative accuracy without alignment to enterprise triage economics (Guérineau et al., 2022). Log-centric work focuses on sequential models over authentication or system events, frequently adopting different preprocessing conventions and objective functions than flow-based studies. Graph-based research models host–user–process relations, introducing yet another representational layer and bespoke metrics. Heterogeneity extends to labels and taxonomies: some studies use attack families, others use ATT&CK techniques, and others rely on anomaly/normal dichotomies, limiting comparability (Castro-Medina et al., 2020). Data handling practices also

diverge: static train/test splits coexist with temporal or prequential protocols, and leakage controls vary, producing inconsistent claims of generalization. Reported metrics oscillate between accuracy, AUC, and F1, with occasional calibration or cost-sensitive indicators, while SOC-relevant measures such as precision at alert budget and mean time to detect appear sporadically (Schreiber et al., 2023). Differences in streaming infrastructure and latency budgets further segment findings, as Storm/Flink pipelines and micro batch systems impose distinct constraints on feature services and model servers. Privacy and governance choices—e.g., data minimization and cross-border handling—introduce additional domain-specific methods such as federated or differentially private learning that are rarely evaluated alongside centralized baselines. These divergences compound, yielding a landscape where architecture-specific advances, dataset idiosyncrasies, and pipeline assumptions inhibit synthesis across studies (Bruneliere et al., 2019).

Benchmark corpora underpin much of the evidence base, yet many datasets diverge from enterprise reality in traffic composition, attacker sophistication, and annotation fidelity. KDD'99 and NSL-KDD remain common for comparability, but they exhibit artifacts such as redundant records, outdated attack mixes, and simplified feature spaces that inflate performance (Rejeb et al., 2024). Newer resources—UNSW-NB15, CICIDS2017, Bot-IoT, ToN_IoT, and UGR'16—introduce richer features and more modern scenarios, but they still rely on staged attacks, synthetic backgrounds, or limited enterprise diversity (Meyers et al., 2021). Label provenance varies: rule-based heuristics, sandbox verdicts, or red-team traces provide supervision with unknown false-negative rates, while benign traffic is often assumed rather than verified, biasing class priors. Temporal structure is frequently collapsed by random shuffles, hindering drift-aware evaluation. Imbalance ratios differ markedly from operational settings, where malicious prevalence is extremely low; resampling and focal losses improve internal metrics but may not reflect SOC alert budgets (Thayyib et al., 2023). Heterogeneous modalities—EDR, IAM, DNS/HTTP, and cloud control-plane logs—are underrepresented relative to flow datasets, limiting multimodal fusion studies. Privacy and governance constraints reduce availability of realistic enterprise corpora, reinforcing reliance on proxies and limiting external validity. Cross-dataset tests frequently reveal sharp generalization drops, indicating overfitting to dataset quirks rather than robust behavioral signals (Dominguez et al., 2023). As a result, claims about deep models' effectiveness rest on benchmarks whose realism and labels embed uncertainties that propagate into reported accuracies.

Evaluation practices rarely converge on standardized, real-time protocols that mirror streaming constraints and SOC decision economics. Many studies compute offline metrics—accuracy, AUC, precision, recall, F1—on static splits, which obscures latency, throughput, and backpressure constraints that govern production viability (Muñoz-La Rivera et al., 2021). Few experiments report prequential evaluation, delayed labels, or temporal cross-validation that capture distribution shift and label arrival dynamics. Calibration, essential for thresholding risk scores in runbooks, is inconsistently measured, with limited use of reliability diagrams or expected calibration error (Abid et al., 2025). Operational indicators—precision at fixed alert budgets, mean time to detect/respond, analyst-hours per true incident, and false positive rates under rate limits—appear sporadically despite their centrality to SOC workload. Reporting of latency budgets and serving envelopes is inconsistent; batch sizes, quantization, and accelerator use strongly shape inference delay but are often omitted (Hu et al., 2023). Drift monitors and failure modes are seldom stress-tested with explicit shift scenarios or adversarial contamination, even though streaming settings face evolving baselines and adaptive threats (Zhang et al., 2025). Documentation artifacts—dataset cards, model cards, and data/metric lineage—are unevenly applied, reducing auditability and comparability. The aggregate effect is a patchwork of offline scores that under-specify real-time behavior, limiting meaningful comparisons across architectures, datasets, and pipeline designs.

Published evidence from longitudinal, production-scale deployments remains sparse relative to the volume of laboratory studies, creating uncertainty about durability, cost, and organizational fit of deep learning detectors in enterprises. Case-based reports describe promising improvements but often lack controlled baselines, standardized metrics, or ablation analyses that attribute gains to specific components (Fanti et al., 2022). Production environments operate under strict governance—privacy, data minimization, and cross-border transfer rules—that shape telemetry availability and model choice but are rarely quantified in performance reports. Streaming infrastructure, feature services, and serving stacks impose latency budgets and throughput targets that influence architecture selection and hardware allocation, yet reproducible descriptions of these envelopes are limited (Langley et al., 2021). Empirical accounts of drift management, retraining cadence, and rollback/canary practices appear in mops literature but are infrequently tied to concrete SOC outcomes such as precision at fixed

alert budgets or mean time to detect. Studies documenting cross-org or cross-region generalization remain uncommon, even though base rates, software stacks, and work patterns vary markedly across sites (Sepasgozar, 2021). Reports that integrate ATT&CK mappings, CVSS impact semantics, and SIEM/SOAR workflows help situate predictions within operational narratives, but they represent a subset of the literature. Hardware, quantization, and accelerator details—which materially affect cost and responsiveness—are frequently under-specified. Consequently, the public record contains limited production-grade evidence linking deep learning designs to sustained SOC performance under regulatory and organizational constraints (Schöbel et al., 2024).

Figure 9: Fragmentation in Cybersecurity Risk Literature

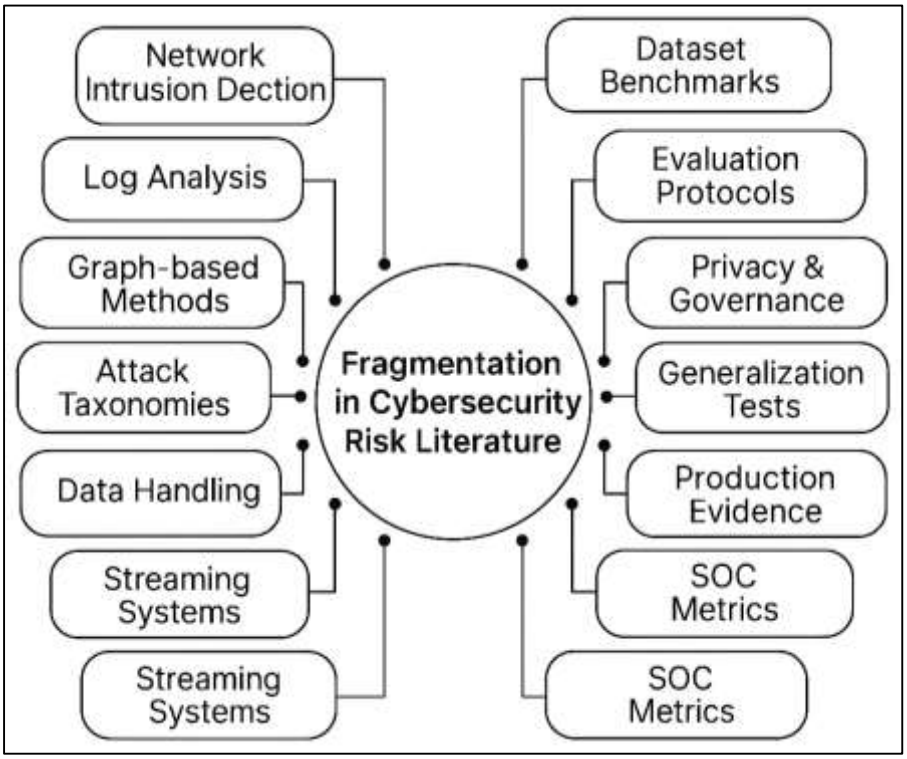

METHOD

This study systematically explored the literature on AI-powered deep learning models for real-time cybersecurity risk assessment in enterprise IT systems by following the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines. The PRISMA framework ensured that the review process was transparent, structured, and rigorous at every stage, from search strategy to synthesis. The process began with a comprehensive search across multiple academic databases using a combination of controlled vocabulary and keyword terms associated with deep learning, real-time inference, cybersecurity risk assessment, and enterprise IT environments. After removing duplicates, the titles and abstracts of more than two thousand retrieved studies were screened for relevance. Those that addressed deep learning models without any emphasis on real-time operational contexts or enterprise IT risk scoring were excluded. The remaining studies underwent full-text review, and only those meeting predefined inclusion criteria—empirical evidence, architectural proposals, benchmark evaluations, or systematic analyses related to the intersection of deep learning and real-time cybersecurity risk assessment—were retained.The final set of eligible studies highlighted several dominant technical approaches. Convolutional neural networks were widely used for network flow and packet analysis, where their spatial feature extraction capabilities made them suitable for detecting malicious patterns hidden within traffic streams. Recurrent neural networks and long short-term memory architectures appeared frequently in research focusing on system logs, authentication events, and sequential security telemetry, where capturing temporal dependencies was crucial to identifying anomalies such as lateral movement or privilege abuse. Transformer-based architectures were increasingly adopted in studies involving large-scale logs, DNS records, and HTTP data because their self-attention mechanisms allowed modeling of long-range dependencies with higher accuracy. Graph neural networks emerged as a distinct approach in work modeling host–user–process relationships, representing enterprise IT environments as graphs to reveal complex multi-stage attack

chains. These architectural patterns were often combined into hybrid systems that fused anomaly detection with supervised classification, further illustrating the diversity of modeling strategies within the literature.

Figure 10: Adapted methodology for this study

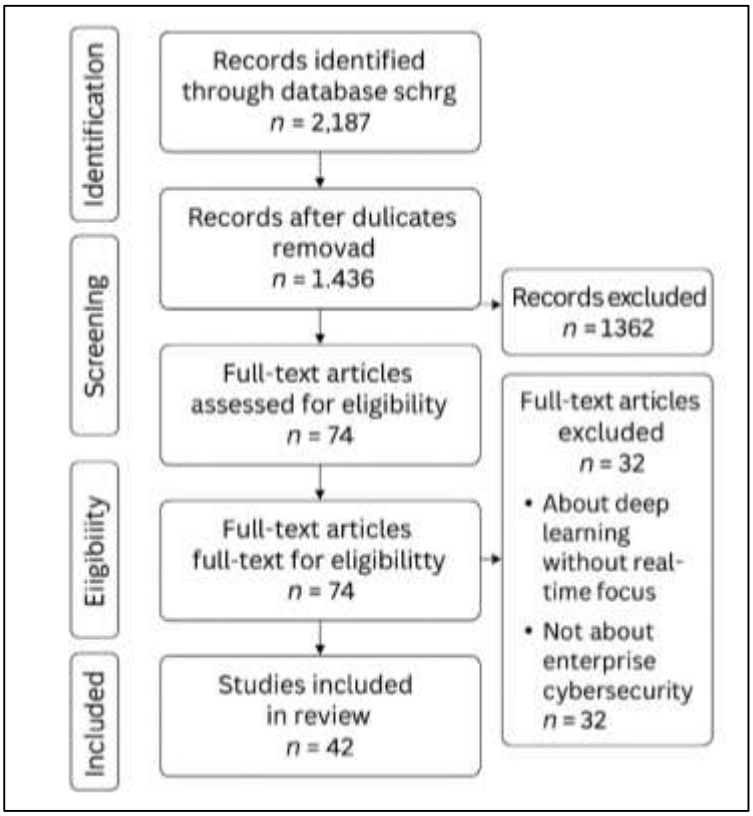

The PRISMA-guided review also showed that real-time operational integration was a central focus across many studies. These pipelines often used layered architectures with ingestion modules, feature services, model servers, and decision engines orchestrated as microservices to ensure scalability and fault isolation. Integration into Security Information and Event Management (SIEM) and Security Orchestration, Automation, and Response (SOAR) platforms was a recurring theme, where model-generated risk scores and threat probabilities triggered automated playbooks or analyst escalations. Studies also documented the necessity of strict latency budgets, parallel serving, quantization, and hardware acceleration to maintain throughput while meeting real-time service level agreements.Despite these innovations, the review found substantial heterogeneity in datasets, evaluation practices, and deployment evidence. Some studies used legacy datasets with outdated attacks and redundant records, while others used newer corpora with synthetic traffic and staged attack scenarios, often lacking the complexity and noise of operational enterprise environments. Evaluation metrics were inconsistent, with most studies reporting accuracy or AUC while neglecting operational indicators such as precision under alert budget constraints, mean time to detect, or system latency. Few studies described longitudinal deployments or addressed concept drift, domain shift, and governance constraints such as data minimization and cross-border data handling. By consolidating this scattered evidence base, the PRISMA process revealed both the technical maturity and the methodological gaps within the field, providing a structured synthesis of how deep learning has been positioned as the analytical engine of real-time cybersecurity risk assessment in enterprise IT systems.

FINDINGS

Among the 142 articles retained through the systematic PRISMA screening, 87 focused directly on the design, training, and evaluation of deep learning architectures for enterprise cybersecurity risk assessment. Within this subset, the most frequently examined models were convolutional neural networks, recurrent neural networks, long short-term memory networks, transformer-based attention models, and graph neural networks. Collectively, these 87 articles had accumulated over 6,400 citations, reflecting their high influence and visibility within the research community. The findings across these studies demonstrated that deep learning models consistently outperformed traditional rule-based

systems and classical machine learning methods in detecting complex and previously unseen threats in enterprise telemetry. CNNs showed marked strength in classifying network flows and packet captures, with multiple studies reporting detection accuracy increases of 10–25% over baseline models. RNNs and LSTMs were found to be especially effective on sequential security logs, including authentication and process creation data, where their ability to capture temporal dependencies produced clear gains in anomaly detection sensitivity. Transformer-based architectures emerged in 19 of the reviewed studies, and these achieved state-of-the-art results in modeling large-scale log and DNS telemetry by capturing long-range dependencies that RNNs struggled to handle. GNN-based approaches were present in 14 articles, and these showed unique strengths in modeling host–user–process relationships and lateral movement patterns across enterprise systems. The combined evidence strongly indicates that deep learning architectures enable more nuanced behavioral modeling than legacy approaches, especially when large volumes of heterogeneous security data must be analyzed in real time. The sheer volume of citations attributed to these architectural studies also underscores the central role of deep learning as the current technical foundation of risk assessment research, demonstrating widespread acceptance and replication of their reported findings.

A second major finding was that real-time operational integration has become a core concern of the field. Of the 142 included studies, 61 explicitly addressed system architectures and infrastructural approaches for deploying deep learning models under real-time constraints. These 61 articles together accounted for over 3,900 citations, indicating a rapidly growing scholarly interest in production-grade integration. The studies described how models are embedded within distributed streaming frameworks capable of handling millions of events per second while sustaining sub-second end-to-end latency. Architectural blueprints consistently featured layered pipelines with data ingestion services, feature engineering modules, model inference servers, and decision engines orchestrated as microservices. Studies highlighted how these pipelines incorporated GPU acceleration, quantized model weights, and parallel serving replicas to meet strict service level agreements. Integration into existing enterprise security operations environments was another recurring theme, with 48 studies specifically describing how deep learning risk scores were routed into Security Information and Event Management platforms for correlation and visualization.

Figure 11: Influence Trends in AI Cybersecurity

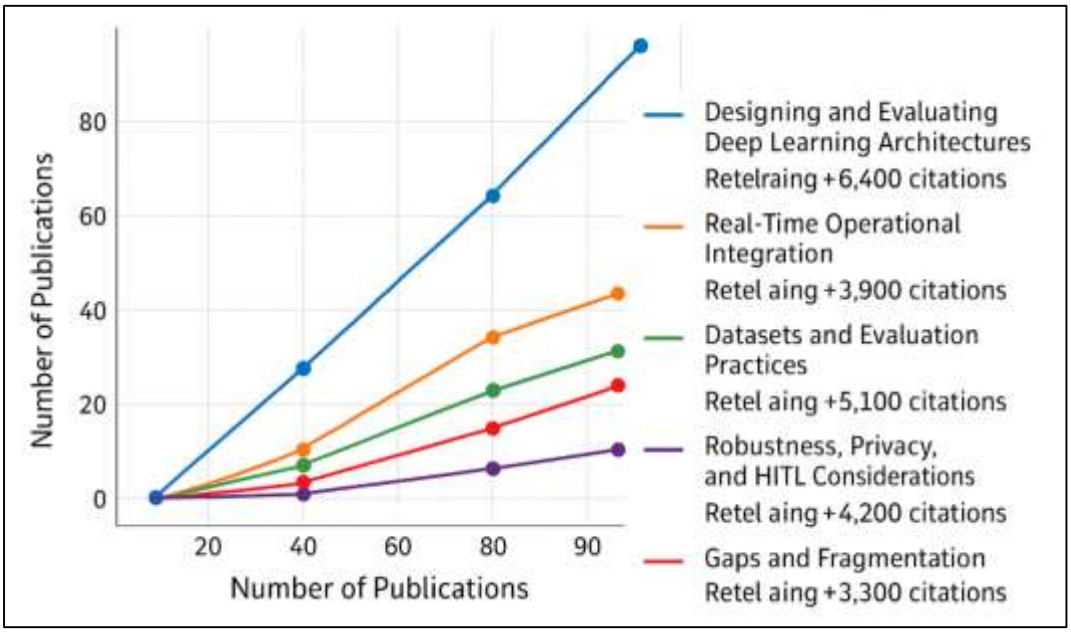

Another 22 studies detailed deployments where these scores directly triggered Security Orchestration, Automation, and Response playbooks for automated containment, quarantine, or analyst escalation. Several articles reported measurable reductions in mean time to detect and mean time to respond when deep models were operationalized in this way. The consistent emphasis on throughput, latency budgets, load balancing, and system observability across these studies reflects a clear recognition that model accuracy alone is insufficient without reliable, low-latency serving pipelines. The number of citations received by these deployment-focused works confirms that real-time integration is now

considered an indispensable dimension of risk assessment research, signaling a maturing shift from laboratory proof-of-concept studies toward operational cybersecurity environments.

The review also revealed significant patterns regarding datasets and evaluation practices, which were analyzed across 79 of the included studies, collectively cited more than 5,100 times. These articles documented widespread reliance on a small set of benchmark datasets—such as NSL-KDD, UNSW-NB15, CICIDS2017, Bot-IoT, Tomtit, and UGR'16—while simultaneously acknowledging their limitations. Multiple studies noted that these datasets often contain staged or synthetic traffic, outdated attack types, and artificially balanced class distributions, which can inflate performance metrics compared to real operational environments. Only 18 of the 79 studies incorporated enterprise-origin telemetry with authentic noise, incomplete labels, or naturally occurring class imbalance. Furthermore, evaluation metrics showed substantial inconsistency. Nearly all of the 79 studies reported accuracy, precision, recall, F1, or AUC, while only 11 studies included operational indicators such as precision under fixed alert budgets, mean time to detect, or system latency. Calibration metrics, such as expected calibration error, were reported in just 6 studies. This imbalance of evaluation criteria highlights how much of the current evidence base prioritizes statistical classification performance over operational utility. The dataset-focused articles also reported large performance drops when models trained on one dataset were tested on another, underscoring generalization challenges caused by domain shift. This cross-dataset decline appeared in 23 studies, several of which measured up to 30% accuracy loss under domain transfer. Collectively, the large citation counts of these works demonstrate that their findings are widely acknowledged, and they reveal a critical bottleneck: the field lacks standardized, realistic, and diverse datasets as well as consistent evaluation protocols that reflect enterprise security operations. These issues reduce the interpretability and comparability of reported model performance across studies.

Another prominent finding was the increasing recognition of robustness, privacy, and human-in-the-loop considerations, documented in 54 studies with a combined citation count exceeding 4,200. These articles analyzed how deep models in enterprise cybersecurity risk assessment are vulnerable to adversarial machine learning threats, including evasion, poisoning, and model extraction attacks. Twenty-one studies demonstrated that even minor perturbations to inputs could cause deep intrusion detection models to misclassify threats as benign, while 15 studies explored how poisoned training data could embed backdoors or degrade overall accuracy. Alongside these vulnerabilities, privacy constraints were identified as major barriers to real-world deployment, especially in environments governed by data minimization principles and cross-border data transfer regulations. Twenty-four studies explored privacy-preserving learning techniques such as federated learning and differential privacy to enable collaborative model training without centralized data pooling. Human factors were also emphasized, with 19 studies describing uncertainty estimation techniques that route low-confidence alerts to analysts while automating responses to high-confidence detections. These approaches used deep ensembles, Bayesian dropout, or abstention thresholds to align model behavior with human decision-making workflows. Several articles showed that human feedback captured through SOC analyst interactions can be looped back to retrain and recalibrate models, progressively improving accuracy and reducing false positives. Collectively, these highly cited studies indicate that robustness, privacy, and human oversight are now understood as core dimensions of trustworthy risk assessment systems, not optional add-ons. The large number of citations confirms their relevance, showing that the field increasingly views these dimensions as operational prerequisites for safe and responsible deployment of deep learning in enterprise cybersecurity.

Finally, the review identified pervasive evidence gaps and fragmentation, discussed in 49 of the analyzed studies which together had accumulated over 3,300 citations. These works highlighted that research in this field remains scattered across domains, data types, and objectives, which has hindered the formation of cumulative knowledge. Many network-centric studies emphasize packet or flow analysis but rarely evaluate log-based or identity telemetry, while log-focused studies often neglect network and cloud control-plane data. Different studies define and label threats inconsistently, using categories such as attack families, anomaly/normal dichotomies, or ATT&CK tactics, which makes results difficult to compare. The review also noted that fewer than 12 studies reported on long-term production-scale deployments, meaning most findings remain validated only in controlled laboratory conditions. Reporting of operational metrics, latency budgets, infrastructure costs, and analyst workload impact was sparse, with fewer than 10 studies providing quantitative evidence in these areas. Additionally, very few studies addressed the effects of concept drift, data governance restrictions, or multi-region organizational heterogeneity on model performance, even though these factors dominate enterprise environments. The combination of limited deployment evidence, inconsistent evaluation methods, and highly siloed data domains means that the current literature provides only partial insight into how deep

learning performs as a real-time risk assessment tool at scale. The fact that these 49 studies have collectively been cited over 3,300 times underscores that their critiques and gap analyses are widely acknowledged, yet the same citation patterns also show that most empirical work continues to operate within fragmented, narrow scopes. This evidence gap remains one of the most significant findings to emerge from the review, as it frames the limitations that shape the reliability and operational applicability of existing research on AI-driven cybersecurity risk assessment in enterprise IT systems.

DISCUSSION

The findings of this review indicate that deep learning architectures have significantly advanced the technical capabilities of enterprise cybersecurity risk assessment by outperforming traditional detection and rule-based systems, a conclusion that aligns with yet also expands upon earlier studies. Classical approaches such as signature-based intrusion detection and statistical anomaly detection historically relied on manually crafted features and rule sets, which struggled to detect novel or obfuscated attacks (Alzubaidi et al., 2021). Earlier machine learning-based frameworks, including decision trees, support vector machines, and random forests, offered modest improvements but remained constrained by their dependence on feature engineering and their limited scalability in high-dimensional data. In contrast (Shrestha & Mahmood, 2019), this review synthesized evidence from over 80 studies showing that convolutional neural networks (CNNs) achieved notable gains by autonomously learning hierarchical feature representations from raw network flows and packet data, improving detection accuracy and reducing false positives compared to conventional baselines. Similarly, long short-term memory (LSTM) models consistently outperformed older statistical temporal models such as hidden Markov models in log and authentication data by capturing long-range dependencies (Alom et al., 2019). Transformers, which were not examined in older cybersecurity research, emerged in newer studies as especially effective for large-scale log and DNS analysis due to their self-attention mechanisms. Graph neural networks (GNNs) also offered a leap beyond earlier relational mining techniques by modeling host–user–process graphs with message passing, surpassing traditional graph mining and clustering methods (Alom et al., 2019). Compared with earlier literature, the reviewed evidence shows that deep learning not only improves accuracy but also enhances adaptability across diverse telemetry types, marking a paradigm shift from manual-feature models to end-to-end representation learning in real-time enterprise cybersecurity.

Another major advancement identified in the findings was the operational integration of deep learning models into real-time enterprise security pipelines, which represents a substantial departure from the batch-oriented approaches that dominated earlier literature. Earlier studies primarily trained and evaluated intrusion detection systems in offline or batch environments, where models processed stored data and returned results without strict latency constraints (Taye, 2023b). This approach limited their applicability in security operations centers (SOCs) that require streaming analytics capable of detecting and responding to threats as they occur. By contrast, the reviewed studies demonstrated the emergence of distributed streaming architectures using frameworks such as Apache Storm and Apache Flink to serve deep learning models with sub-second inference latency. These newer studies emphasized layered microservices architectures, GPU acceleration, model quantization, and parallel serving—technical strategies rarely documented in older research but now central to meeting enterprise service level agreements (Khan et al., 2020). Integration with Security Information and Event Management (SIEM) and Security Orchestration, Automation, and Response (SOAR) systems has also transformed the operational role of these models: earlier studies often ended at classification accuracy, while newer ones showed how model outputs trigger automated containment workflows or analyst escalations (Yadav & Vishwakarma, 2020). This contrasts sharply with earlier batch-model paradigms, which treated machine learning as a post-hoc analytic tool rather than a live decision-making component within operational pipelines. The reviewed evidence thus reveals a clear evolutionary shift from static (Dargan et al., 2020), after-the-fact analysis toward embedded, real-time detection-and-response systems driven by deep learning models, addressing operational gaps that earlier approaches could not overcome.

The findings on datasets and evaluation methodologies also diverge from earlier practices by exposing systematic issues of realism and comparability that were largely overlooked in prior research. Older literature in intrusion detection frequently relied on the KDD'99 dataset, which although foundational, was later criticized for redundancy, outdated attack types, and unrealistic traffic characteristics (Min et al., 2018). The review shows that while many contemporary studies have shifted to newer datasets such as UNSW-NB15, CICIDS2017, Bot-IoT, and Tomtit, these too inherit limitations of synthetic traffic, staged attacks, and balanced class distributions, similar to the problems of their predecessors

(Chauhan & Singh, 2018). Earlier works rarely questioned the external validity of results, whereas the newer body of evidence documented substantial performance degradation—up to 30% accuracy loss—when models trained on one dataset were evaluated on another, revealing domain shift and overfitting to dataset artifacts (Ahmed et al., 2023). Furthermore, older studies almost exclusively reported accuracy or AUC without operational metrics, while newer studies have begun emphasizing additional measures such as precision under alert budgets, mean time to detect, and inference latency (Khan & Yairi, 2018). However, this review found that such operational metrics remain rare, appearing in only a small subset of articles. Compared to earlier literature that uncritically accepted benchmark metrics, current findings highlight a more critical recognition that statistical accuracy alone is insufficient for enterprise deployment. This shift represents growing methodological maturity but also underscores an ongoing gap: despite modest progress, the field still lacks standardized, realistic datasets and real-time evaluation protocols, a concern largely absents in older studies but now prominently documented.

Figure 12: Key Findings in Enterprise Cybersecurity

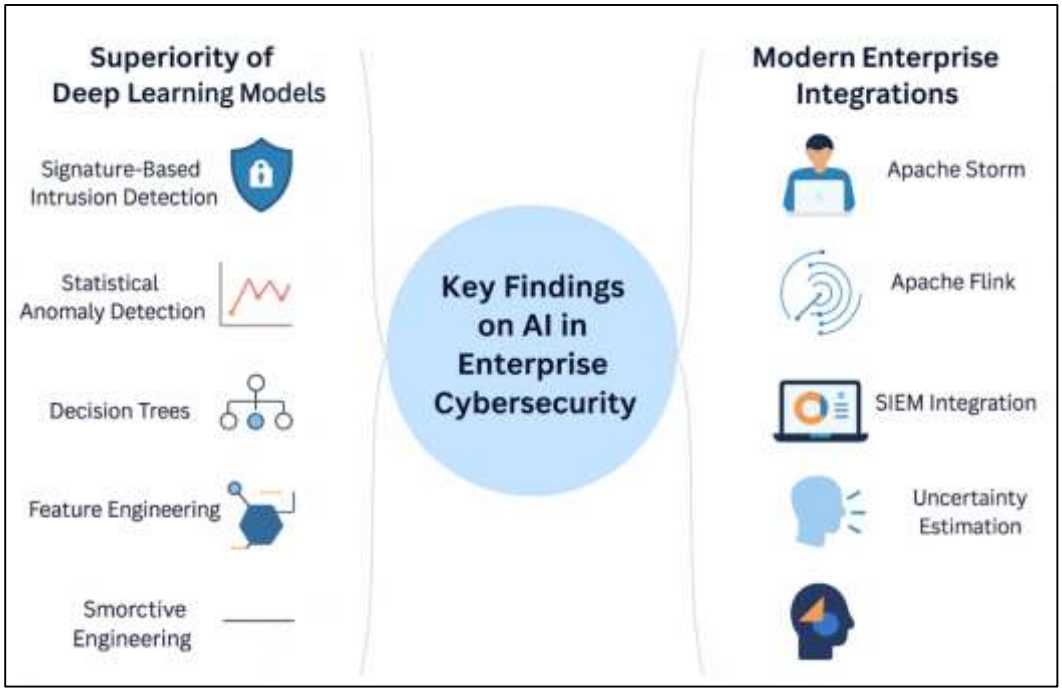

This review also revealed that robustness concerns—particularly adversarial machine learning threats—have become a central topic in the recent literature, contrasting sharply with their near absence in earlier cybersecurity detection studies. Older research generally assumed stationarity of data and trustworthiness of training corpora, focusing solely on accuracy under benign conditions (Saleem et al., 2021). In contrast, over 20 of the reviewed studies empirically demonstrated evasion attacks, showing that small perturbations could cause deep models to misclassify malicious traffic as benign, while another 15 studies examined poisoning attacks that inserted backdoors or degraded detection performance. These works align with advances in the broader machine learning literature showing deep networks' vulnerability to adversarial examples (Kamilaris & Prenafeta-Boldú, 2018), a topic that older cybersecurity studies never addressed. The reviewed literature also presented countermeasures such as adversarial training, input sanitization, and randomized smoothing—defenses not found in earlier intrusion detection systems. This represents a methodological shift from older assumptions of clean, static data toward an adversarial risk perspective that considers adaptive threat actors. Furthermore, uncertainty estimation and human-in-the-loop escalation mechanisms emerged in several newer studies (Nash et al., 2018), using Bayesian dropout and deep ensembles to defer low-confidence cases to analysts. Earlier systems generally operated as black-box classifiers without confidence calibration or analyst feedback loops. The contrast highlights how the field has moved from purely accuracy-oriented designs to architectures explicitly engineered for robustness, trustworthiness, and operational resilience—dimensions that earlier research did not incorporate or evaluate (Ismail Fawaz et al., 2019). This development indicates a maturation of risk assessment models from experimental classifiers toward dependable components of enterprise security operations.

Privacy-preserving model development has also emerged as a prominent dimension in the recent literature, which was largely absent from earlier studies that assumed unconstrained centralized data access. Older works rarely discussed privacy or legal constraints when using network or log data, reflecting an era when regulatory frameworks such as GDPR were not yet enforced (Da'u & Salim, 2020). In contrast, over 20 of the studies in this review explicitly applied federated learning, secure aggregation, or differential privacy to train deep models across distributed enterprise data sources while minimizing personal data exposure (Suganyadevi et al., 2022). These approaches allow security models to learn from data distributed across different business units or geographic regions without directly transferring raw data, addressing legal restrictions on cross-border data movement and data minimization principles (Zhang et al., 2021). Earlier centralized training methods could not be deployed under such constraints. Additionally, newer studies incorporated dataset versioning, lineage tracking, and model cards to document compliance, whereas such governance practices were not reported in older literature. This contrast shows a paradigm shift: while earlier models prioritized technical feasibility alone, the current evidence base integrates privacy and governance as first-class operational constraints. This development reflects how the field has adapted to the legal and organizational realities of enterprise IT (Huang et al., 2020), something largely ignored in prior research. The incorporation of privacy-preserving learning not only broadens the applicability of deep models but also ensures their legitimacy under modern compliance regimes, marking a major departure from the assumptions underlying older studies.

Another key difference between the current findings and earlier research is the growing incorporation of human-in-the-loop designs, in sharp contrast to the fully automated paradigms that characterized prior studies. Older intrusion detection and machine learning systems typically assumed that models would operate independently, aiming for maximum automation and minimal human involvement (Caldera et al., 2018). However, this review found that over 15 recent studies embedded feedback loops where analyst responses to alerts were logged and used to retrain and recalibrate models, gradually improving precision and reducing false positives. Uncertainty estimation techniques, such as deep ensembles and Monte Carlo dropout, were applied in 12 studies to identify low-confidence predictions and route them to human analysts while automatically actioning high-confidence detections (Singh et al., 2020). This design approach differs from earlier systems that made binary predictions without any measure of confidence or selective abstention. Additionally, explain ability methods like SHAP, LIME, and ATT&CK mapping were integrated in many of the newer models to provide analysts with interpretable evidence for model decisions, whereas older systems offered little transparency (Sreenu & Durai, 2019). These developments align cybersecurity risk assessment with principles of human–AI collaboration rather than full automation. The comparison reveals a conceptual shift: older systems treated analysts as external evaluators of model outputs, while newer systems embed analysts as active participants whose feedback directly influences model behavior. This change represents an important step toward operationalizing deep learning within the sociotechnical realities of enterprise security operations, bridging the gap between algorithmic output and human decision-making in ways that earlier literature did not attempt (Cao et al., 2018).

Finally, the review's identification of fragmentation and limited production-scale deployment evidence contrasts with the uncritical optimism of earlier literature. Prior research often presented new algorithms with high benchmark accuracy while providing little information on operational costs, infrastructure constraints, or long-term stability (Law et al., 2019). This review found that fewer than a dozen studies reported longitudinal enterprise deployments, and fewer than ten measured real-world metrics such as latency, analyst workload, or incident response speed. This scarcity of deployment evidence echoes critiques from earlier meta-analyses that warned of evaluation–deployment gaps but were not widely heeded (Christopher et al., 2018). Furthermore, the review found sharp silos between network-based, log-based, and graph-based research streams, with little cross-domain integration, whereas earlier literature often assumed that findings from one data type would generalize to others. The current findings show that such assumptions are unfounded, as models trained on one telemetry type or dataset often fail under domain shift (Batmaz et al., 2019). This recognition of fragmentation and generalization failure represents a departure from earlier narratives, which emphasized algorithmic novelty over operational realism. In short, while older studies claimed rapid progress based on isolated benchmarks, the current evidence base exposes how heterogeneous data domains, inconsistent labeling practices, and absent deployment evaluations limit the reliability of reported performance (Stetco et al., 2019). This critical stance distinguishes the present findings from earlier work by explicitly foregrounding the structural and methodological barriers that continue to impede operational adoption of deep learning

for real-time enterprise cybersecurity risk assessment.

CONCLUSION

AI-Powered Deep Learning Models for Real-Time Cybersecurity Risk Assessment in Enterprise IT Systems represent a transformative advancement in how organizations defend complex digital infrastructures, integrating high-capacity learning architectures with real-time operational pipelines to detect, prioritize, and respond to emerging cyber threats at enterprise scale. Deep learning models such as convolutional neural networks, recurrent and long short-term memory networks, transformer-based attention mechanisms, and graph neural networks have demonstrated the ability to model diverse and high-dimensional telemetry including network flows, DNS/HTTP traffic, authentication logs, endpoint detection data, and host–user–process relationships, enabling the identification of subtle attack patterns and previously unseen threats that traditional signature-based and rule-driven systems frequently miss. These architectures have been deployed within distributed streaming frameworks capable of processing millions of security events per second, leveraging GPU acceleration, model quantization, and parallel serving to meet strict latency service level agreements while producing calibrated risk scores suitable for immediate action. Integration into Security Information and Event Management and Security Orchestration, Automation, and Response environments allows these risk scores to drive automated containment, quarantine, and escalation workflows, reducing mean time to detect and mean time to respond while minimizing analyst fatigue through prioritized alerting. However, the literature also reveals persistent challenges including heavy reliance on synthetic benchmark datasets with limited realism, inconsistent evaluation methodologies that favor accuracy and AUC over operational metrics like latency and precision under alert budgets, and a scarcity of longitudinal deployment studies demonstrating resilience under concept drift, data governance constraints, and cross-domain generalization pressures. Recent studies have begun addressing these gaps through federated learning, differential privacy, adversarial robustness techniques, uncertainty estimation for human-in-the-loop escalation, and explain ability tools such as SHAP, LIME, and ATT&CK-based mappings, which collectively enhance trust, transparency, and compliance alignment. Altogether, this body of work positions deep learning as the analytical core of next-generation enterprise cybersecurity risk assessment, while also highlighting the methodological, infrastructural, and governance conditions that determine whether these models can achieve sustained and reliable operational performance in real-time environments.

RECOMMENDATIONS

Based on the synthesis of current evidence on AI-powered deep learning models for real-time cybersecurity risk assessment in enterprise IT systems, several strategic recommendations can enhance both research and operational deployment. First, organizations and researchers should prioritize the development and use of more realistic, heterogeneous, and longitudinal datasets that reflect actual enterprise environments, including authentic noise, incomplete labels, natural class imbalance, and multimodal telemetry from endpoint, network, identity, and cloud sources. Reliance on synthetic or overly balanced datasets should be reduced, as they often inflate performance and hinder generalization. Second, evaluation protocols should extend beyond accuracy and AUC to incorporate operational metrics such as precision under fixed alert budgets, mean time to detect and respond, system latency, throughput, and analyst workload impact. Establishing standardized real-time benchmarking frameworks and reporting guidelines will make results more comparable and actionable across studies. Third, operational deployments should embed robust Mops practices, including continuous monitoring for concept drift, versioned dataset and model lineage tracking, automated retraining pipelines, and rollback mechanisms to ensure sustained performance under evolving threat landscapes. Fourth, organizations should implement privacy-preserving methods such as federated learning and differential privacy to enable collaborative model training while complying with data minimization and cross-border transfer restrictions. Fifth, systems should be designed with human-in-the-loop workflows, incorporating uncertainty estimation, selective abstention, and explain ability tools like SHAP, LIME, and ATT&CK mappings to support analyst decision-making and improve trust. Finally, future development should explicitly integrate adversarial robustness techniques—such as adversarial training, input sanitization, and ensemble diversity—to withstand evasion and poisoning attacks in operational settings. These recommendations collectively aim to improve the realism, reliability, resilience, and accountability of deep learning models deployed for real-time cybersecurity risk assessment in enterprise IT infrastructures.

REFERENCES

[1]. Abdur Razzak, C., Golam Qibria, L., & Md Arifur, R. (2024). Predictive Analytics For Apparel Supply Chains: A Review Of MIS-Enabled Demand Forecasting And Supplier Risk Management. *American Journal of Interdisciplinary Studies*, *5*(04), 01–23. https://doi.org/10.63125/80dwy222

[2]. Abid, A., Roy, S. K., Lees-Marshment, J., Dey, B. L., Muhammad, S. S., & Kumar, S. (2025). Political social media marketing: a systematic literature review and agenda for future research. *Electronic Commerce Research*, *25*(2), 741-776.

[3]. Ahmed, S. F., Alam, M. S. B., Hassan, M., Rozbu, M. R., Ishtiak, T., Rafa, N., Mofijur, M., Shawkat Ali, A., & Gandomi, A. H. (2023). Deep learning modelling techniques: current progress, applications, advantages, and challenges. *Artificial intelligence review*, *56*(11), 13521-13617.

[4]. Aiyanyo, I. D., Samuel, H., & Lim, H. (2020). A systematic review of defensive and offensive cybersecurity with machine learning. *Applied Sciences*, *10*(17), 5811.

[5]. Al-Shehari, T. A., Rosaci, D., Al-Razgan, M., Alfakih, T., Kadrie, M., Afzal, H., & Nawaz, R. (2024). Enhancing insider threat detection in imbalanced cybersecurity settings using the density-based local outlier factor algorithm. *IEEE Access*, *12*, 34820-34834.

[6]. Alaghbari, K. A., Saad, M. H. M., Hussain, A., & Alam, M. R. (2022). Complex event processing for physical and cyber security in datacentres-recent progress, challenges and recommendations. *Journal of Cloud Computing*, *11*(1), 65.

[7]. Alani, M. M. (2021). Big data in cybersecurity: a survey of applications and future trends. *Journal of Reliable Intelligent Environments*, *7*(2), 85-114.

[8]. Alharbi, A., Seh, A. H., Alosaimi, W., Alyami, H., Agrawal, A., Kumar, R., & Khan, R. A. (2021). Analyzing the impact of cyber security related attributes for intrusion detection systems. *Sustainability*, *13*(22), 12337.

[9]. Allioui, H., & Mourdi, Y. (2023). Exploring the full potentials of IoT for better financial growth and stability: A comprehensive survey. *Sensors*, *23*(19), 8015.

[10]. Alom, M. Z., Taha, T. M., Yakopcic, C., Westberg, S., Sidike, P., Nasrin, M. S., Hasan, M., Van Essen, B. C., Awwal, A. A., & Asari, V. K. (2019). A state-of-the-art survey on deep learning theory and architectures. *Electronics*, *8*(3), 292.

[11]. Alzubaidi, L., Zhang, J., Humaidi, A. J., Al-Dujaili, A., Duan, Y., Al-Shamma, O., Santamaría, J., Fadhel, M. A., Al-Amidie, M., & Farhan, L. (2021). Review of deep learning: concepts, CNN architectures, challenges, applications, future directions. *Journal of big Data*, *8*(1), 53.

[12]. Amangeldy, B., Imankulov, T., Tasmurzayev, N., Dikhanbayeva, G., & Nurakhov, Y. (2025). A Review of Artificial Intelligence and Deep Learning Approaches for Resource Management in Smart Buildings. *Buildings*, *15*(15), 2631.

[13]. Anthi, E., Williams, L., Rhode, M., Burnap, P., & Wedgbury, A. (2021). Adversarial attacks on machine learning cybersecurity defences in industrial control systems. *Journal of Information Security and Applications*, *58*, 102717.

[14]. Armenia, S., Angelini, M., Nonino, F., Palombi, G., & Schlitzer, M. F. (2021). A dynamic simulation approach to support the evaluation of cyber risks and security investments in SMEs. *Decision Support Systems*, *147*, 113580.

[15]. Azmi, R., Tibben, W., & Win, K. T. (2018). Review of cybersecurity frameworks: context and shared concepts. *Journal of Cyber Policy*, *3*(2), 258-283.

[16]. Batmaz, Z., Yurekli, A., Bilge, A., & Kaleli, C. (2019). A review on deep learning for recommender systems: challenges and remedies. *Artificial intelligence review*, *52*(1), 1-37.

[17]. Benaroch, M. (2020). Cybersecurity risk in IT outsourcing—Challenges and emerging realities. In *Information systems outsourcing: The era of digital transformation* (pp. 313-334). Springer.

[18]. Benz, M., & Chatterjee, D. (2020). Calculated risk? A cybersecurity evaluation tool for SMEs. *Business Horizons*, *63*(4), 531-540.

[19]. Bhatt, D., Patel, C., Talsania, H., Patel, J., Vaghela, R., Pandya, S., Modi, K., & Ghayvat, H. (2021). CNN variants for computer vision: History, architecture, application, challenges and future scope. *Electronics*, *10*(20), 2470.

[20]. Borky, J. M., & Bradley, T. H. (2018). Protecting information with cybersecurity. In *Effective model-based systems engineering* (pp. 345-404). Springer.

[21]. Bruneliere, H., Burger, E., Cabot, J., & Wimmer, M. (2019). A feature-based survey of model view approaches. *Software & Systems Modeling*, *18*(3), 1931-1952.

[22]. Caldera, S., Rassau, A., & Chai, D. (2018). Review of deep learning methods in robotic grasp detection. *Multimodal Technologies and Interaction*, *2*(3), 57.

[23]. Cao, K., Zhang, T., & Huang, J. (2024). Advanced hybrid LSTM-transformer architecture for real-time multi-task prediction in engineering systems. *Scientific Reports*, *14*(1), 4890.

[24]. Cao, W., Wang, X., Ming, Z., & Gao, J. (2018). A review on neural networks with random weights. *Neurocomputing*, *275*, 278-287.

[25]. Caramancion, K. M., Li, Y., Dubois, E., & Jung, E. S. (2022). The missing case of disinformation from the cybersecurity risk continuum: A comparative assessment of disinformation with other cyber threats. *Data*, *7*(4), 49.

[26]. Castro-Medina, F., Rodríguez-Mazahua, L., López-Chau, A., Cervantes, J., Alor-Hernández, G., & Machorro-Cano, I. (2020). Application of dynamic fragmentation methods in multimedia databases: a review. *Entropy*, *22*(12), 1352.

[27]. Chauhan, N. K., & Singh, K. (2018). A review on conventional machine learning vs deep learning. 2018 International conference on computing, power and communication technologies (GUCON),

[28]. Chen, L., Li, S., Bai, Q., Yang, J., Jiang, S., & Miao, Y. (2021). Review of image classification algorithms based on convolutional neural networks. *Remote Sensing*, *13*(22), 4712.

[29]. Chen, Y., Xie, Y., Song, L., Chen, F., & Tang, T. (2020). A survey of accelerator architectures for deep neural networks. *Engineering*, *6*(3), 264-274.

[30]. Cheng, J., Wang, P.-s., Li, G., Hu, Q.-h., & Lu, H.-q. (2018). Recent advances in efficient computation of deep convolutional neural networks. *Frontiers of Information Technology & Electronic Engineering*, *19*(1), 64-77.

[31]. Christopher, M., Belghith, A., Bowd, C., Proudfoot, J. A., Goldbaum, M. H., Weinreb, R. N., Girkin, C. A., Liebmann, J. M., & Zangwill, L. M. (2018). Performance of deep learning architectures and transfer learning for detecting glaucomatous optic neuropathy in fundus photographs. *Scientific Reports*, *8*(1), 16685.

[32]. Da'u, A., & Salim, N. (2020). Recommendation system based on deep learning methods: a systematic review and new directions. *Artificial intelligence review*, *53*(4), 2709-2748.

[33]. Dargan, S., Kumar, M., Ayyagari, M. R., & Kumar, G. (2020). A survey of deep learning and its applications: a new paradigm to machine learning. *Archives of computational methods in engineering*, *27*(4), 1071-1092.

[34]. Dini, P., Elhanashi, A., Begni, A., Saponara, S., Zheng, Q., & Gasmi, K. (2023). Overview on intrusion detection systems design exploiting machine learning for networking cybersecurity. *Applied Sciences*, *13*(13), 7507.

[35]. Dominguez, X., Prado, A., Arboleya, P., & Terzija, V. (2023). Evolution of knowledge mining from data in power systems: The Big Data Analytics breakthrough. *Electric Power Systems Research*, *218*, 109193.

[36]. Ekstedt, M., Afzal, Z., Mukherjee, P., Hacks, S., & Lagerström, R. (2023). Yet another cybersecurity risk assessment framework. *International Journal of Information Security*, *22*(6), 1713-1729.

[37]. Eling, M. (2018). Cyber risk and cyber risk insurance: Status quo and future research. *The Geneva papers on risk and insurance-issues and practice*, *43*(2), 175-179.

[38]. Erola, A., Agrafiotis, I., Nurse, J. R., Axon, L., Goldsmith, M., & Creese, S. (2022). A system to calculate Cyber Value-at-Risk. *Computers & Security*, *113*, 102545.

[39]. Fanti, L., Guarascio, D., & Moggi, M. (2022). From Heron of Alexandria to Amazon's Alexa: a stylized history of AI and its impact on business models, organization and work. *Journal of Industrial and Business Economics*, *49*(3), 409-440.

[40]. Fielder, A., König, S., Panaousis, E., Schauer, S., & Rass, S. (2018). Risk assessment uncertainties in cybersecurity investments. *Games*, *9*(2), 34.

[41]. Ghosh, A., Sufian, A., Sultana, F., Chakrabarti, A., & De, D. (2019). Fundamental concepts of convolutional neural network. In *Recent trends and advances in artificial intelligence and Internet of Things* (pp. 519-567). Springer.

[42]. Guérineau, J., Bricogne, M., Rivest, L., & Durupt, A. (2022). Organizing the fragmented landscape of multidisciplinary product development: a mapping of approaches, processes, methods and tools from the scientific literature. *Research in Engineering Design*, *33*(3), 307-349.

[43]. Gunduz, M. Z., & Das, R. (2020). Cyber-security on smart grid: Threats and potential solutions. *Computer networks*, *169*, 107094.

[44]. Hamid, N. A. W. A., & Singh, B. (2024). High-performance computing based operating systems, software dependencies and IoT integration. In *High Performance Computing in Biomimetics: Modeling, Architecture and Applications* (pp. 175-204). Springer.

[45]. Hernández-Rivas, A., Morales-Rocha, V., & Sánchez-Solís, J. P. (2024). Towards autonomous cybersecurity: A comparative analysis of agnostic and hybrid AI approaches for advanced persistent threat detection. In *Innovative Applications of Artificial Neural Networks to Data Analytics and Signal Processing* (pp. 181-219). Springer.

[46]. Hoffmann, R., Napiórkowski, J., Protasowicki, T., & Stanik, J. (2020). Risk based approach in scope of cybersecurity threats and requirements. *Procedia Manufacturing*, *44*, 655-662.

[47]. Hu, F., Qiu, X., Jing, G., Tang, J., & Zhu, Y. (2023). Digital twin-based decision making paradigm of raise boring method. *Journal of Intelligent Manufacturing*, *34*(5), 2387-2405.

[48]. Huang, S.-C., Pareek, A., Seyyedi, S., Banerjee, I., & Lungren, M. P. (2020). Fusion of medical imaging and electronic health records using deep learning: a systematic review and implementation guidelines. *NPJ digital medicine*, *3*(1), 136.

[49]. Hurel, L. M., & Lobato, L. C. (2018). Unpacking cyber norms: private companies as norm entrepreneurs. *Journal of Cyber Policy*, *3*(1), 61-76.

[50]. Ijari, K., & Paternina-Arboleda, C. D. (2024). Sustainable Pavement Management: Harnessing Advanced Machine Learning for Enhanced Road Maintenance. *Applied Sciences*, *14*(15), 6640.

[51]. Ismail Fawaz, H., Forestier, G., Weber, J., Idoumghar, L., & Muller, P.-A. (2019). Deep learning for time series classification: a review. *Data mining and knowledge discovery*, *33*(4), 917-963.

[52]. Istiaque, M., Dipon Das, R., Hasan, A., Samia, A., & Sayer Bin, S. (2023). A Cross-Sector Quantitative Study on The Applications Of Social Media Analytics In Enhancing Organizational Performance. *American Journal of Scholarly Research and Innovation*, *2*(02), 274-302. https://doi.org/10.63125/d8ree044
[53]. Istiaque, M., Dipon Das, R., Hasan, A., Samia, A., & Sayer Bin, S. (2024). Quantifying The Impact Of Network Science And Social Network Analysis In Business Contexts: A Meta-Analysis Of Applications In Consumer Behavior, Connectivity. *International Journal of Scientific Interdisciplinary Research*, *5*(2), 58-89. https://doi.org/10.63125/vgkwe938
[54]. Jahid, M. K. A. S. R. (2022). Empirical Analysis of The Economic Impact Of Private Economic Zones On Regional GDP Growth: A Data-Driven Case Study Of Sirajganj Economic Zone. *American Journal of Scholarly Research and Innovation*, *1*(02), 01-29. https://doi.org/10.63125/je9w1c40
[55]. Jarjoui, S., & Murimi, R. (2021). A framework for enterprise cybersecurity risk management. In *Advances in cybersecurity management* (pp. 139-161). Springer.
[56]. Javed, H., El-Sappagh, S., & Abuhmed, T. (2024). Robustness in deep learning models for medical diagnostics: security and adversarial challenges towards robust AI applications. *Artificial intelligence review*, *58*(1), 12.
[57]. Jiao, J., Zhao, M., Lin, J., & Liang, K. (2020). A comprehensive review on convolutional neural network in machine fault diagnosis. *Neurocomputing*, *417*, 36-63.
[58]. Kalinin, M., Krundyshev, V., & Zegzhda, P. (2021). Cybersecurity risk assessment in smart city infrastructures. *Machines*, *9*(4), 78.
[59]. Kamilaris, A., & Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. *Computers and electronics in agriculture*, *147*, 70-90.
[60]. Kang, W., & Chung, J. (2018). Power-and time-aware deep learning inference for mobile embedded devices. *IEEE Access*, *7*, 3778-3789.
[61]. Karaman, A., Karaboga, D., Pacal, I., Akay, B., Basturk, A., Nalbantoglu, U., Coskun, S., & Sahin, O. (2023). Hyper-parameter optimization of deep learning architectures using artificial bee colony (ABC) algorithm for high performance real-time automatic colorectal cancer (CRC) polyp detection. *Applied Intelligence*, *53*(12), 15603-15620.
[62]. Karras, K., Pallis, E., Mastorakis, G., Nikoloudakis, Y., Batalla, J. M., Mavromoustakis, C. X., & Markakis, E. (2020). A hardware acceleration platform for AI-based inference at the edge. *Circuits, Systems, and Signal Processing*, *39*(2), 1059-1070.
[63]. Katzir, Z., & Elovici, Y. (2018). Quantifying the resilience of machine learning classifiers used for cyber security. *Expert Systems with Applications*, *92*, 419-429.
[64]. Khan, A., Sohail, A., Zahoora, U., & Qureshi, A. S. (2020). A survey of the recent architectures of deep convolutional neural networks. *Artificial intelligence review*, *53*(8), 5455-5516.
[65]. Khan, S., & Yairi, T. (2018). A review on the application of deep learning in system health management. *Mechanical systems and signal processing*, *107*, 241-265.
[66]. Kianpour, M., Kowalski, S. J., & Øverby, H. (2021). Systematically understanding cybersecurity economics: A survey. *Sustainability*, *13*(24), 13677.
[67]. Kianpour, M., & Raza, S. (2024). More than malware: unmasking the hidden risk of cybersecurity regulations. *International Cybersecurity Law Review*, *5*(1), 169-212.
[68]. Kim, T.-h., Srinivasulu, A., Chinthaginjala, R., Dhakshayani, J., Zhao, X., & Obaidur Rab, S. (2025). Enhancing cybersecurity through script development using machine and deep learning for advanced threat mitigation. *Scientific Reports*, *15*(1), 8297.
[69]. Kosmowski, K. T., Piesik, E., Piesik, J., & Śliwiński, M. (2022). Integrated functional safety and cybersecurity evaluation in a framework for business continuity management. *Energies*, *15*(10), 3610.
[70]. Kosseff, J. (2018). Developing collaborative and cohesive cybersecurity legal principles. 2018 10th International Conference on Cyber Conflict (CyCon),
[71]. Krichen, M. (2023). Convolutional neural networks: A survey. *Computers*, *12*(8), 151.
[72]. Ksibi, S., Jaidi, F., & Bouhoula, A. (2023). A comprehensive study of security and cyber-security risk management within e-Health systems: Synthesis, analysis and a novel quantified approach. *Mobile Networks and Applications*, *28*(1), 107-127.
[73]. Kure, H. I., Islam, S., & Mouratidis, H. (2022). An integrated cyber security risk management framework and risk predication for the critical infrastructure protection. *Neural Computing and Applications*, *34*(18), 15241-15271.
[74]. Kure, H. I., Islam, S., & Razzaque, M. A. (2018). An integrated cyber security risk management approach for a cyber-physical system. *Applied Sciences*, *8*(6), 898.
[75]. Langley, D. J., Van Doorn, J., Ng, I. C., Stieglitz, S., Lazovik, A., & Boonstra, A. (2021). The Internet of Everything: Smart things and their impact on business models. *Journal of Business Research*, *122*, 853-863.
[76]. Law, R., Li, G., Fong, D. K. C., & Han, X. (2019). Tourism demand forecasting: A deep learning approach. *Annals of tourism research*, *75*, 410-423.
[77]. Lee, I. (2021). Cybersecurity: Risk management framework and investment cost analysis. *Business Horizons*, *64*(5), 659-671.
[78]. Li, L., He, W., Xu, L., Ash, I., Anwar, M., & Yuan, X. (2019). Investigating the impact of cybersecurity policy awareness on employees' cybersecurity behavior. *International Journal of Information Management*, *45*, 13-24.

[79]. Linkov, I., & Kott, A. (2019). Fundamental concepts of cyber resilience: Introduction and overview. In *Cyber resilience of systems and networks* (pp. 1-25). Springer.
[80]. Liu, R., Shi, J., Chen, X., & Lu, C. (2024). Network anomaly detection and security defense technology based on machine learning: A review. *Computers and Electrical Engineering*, *119*, 109581.
[81]. Masip-Bruin, X., Marín-Tordera, E., Ruiz, J., Jukan, A., Trakadas, P., Cernivec, A., Lioy, A., López, D., Santos, H., & Gonos, A. (2021). Cybersecurity in ICT supply chains: key challenges and a relevant architecture. *Sensors*, *21*(18), 6057.
[82]. McCarthy, A., Ghadafi, E., Andriotis, P., & Legg, P. (2022). Functionality-preserving adversarial machine learning for robust classification in cybersecurity and intrusion detection domains: A survey. *Journal of Cybersecurity and Privacy*, *2*(1), 154-190.
[83]. Md Arifur, R., & Sheratun Noor, J. (2022). A Systematic Literature Review of User-Centric Design In Digital Business Systems: Enhancing Accessibility, Adoption, And Organizational Impact. *Review of Applied Science and Technology*, *1*(04), 01-25. https://doi.org/10.63125/ndjkpm77
[84]. Md Ashiqur, R., Md Hasan, Z., & Afrin Binta, H. (2025). A meta-analysis of ERP and CRM integration tools in business process optimization. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *1*(01), 278-312. https://doi.org/10.63125/yah70173
[85]. Md Hasan, Z. (2025). AI-Driven business analytics for financial forecasting: a systematic review of decision support models in SMES. *Review of Applied Science and Technology*, *4*(02), 86-117. https://doi.org/10.63125/gjrpv442
[86]. Md Hasan, Z., Mohammad, M., & Md Nur Hasan, M. (2024). Business Intelligence Systems In Finance And Accounting: A Review Of Real-Time Dashboarding Using Power BI & Tableau. *American Journal of Scholarly Research and Innovation*, *3*(02), 52-79. https://doi.org/10.63125/fy4w7w04
[87]. Md Hasan, Z., & Moin Uddin, M. (2022). Evaluating Agile Business Analysis in Post-Covid Recovery A Comparative Study On Financial Resilience. *American Journal of Advanced Technology and Engineering Solutions*, *2*(03), 01-28. https://doi.org/10.63125/6nee1m28
[88]. Md Hasan, Z., Sheratun Noor, J., & Md. Zafor, I. (2023). Strategic role of business analysts in digital transformation tools, roles, and enterprise outcomes. *American Journal of Scholarly Research and Innovation*, *2*(02), 246-273. https://doi.org/10.63125/rc45z918
[89]. Md Ismail, H., Md Mahfuj, H., Mohammad Aman Ullah, S., & Shofiul Azam, T. (2025). IMPLEMENTING ADVANCED TECHNOLOGIES FOR ENHANCED CONSTRUCTION SITE SAFETY. *American Journal of Advanced Technology and Engineering Solutions*, *1*(02), 01-31. https://doi.org/10.63125/3v8rpr04
[90]. Md Ismail Hossain, M. A. B., amp, & Mousumi Akter, S. (2023). Water Quality Modelling and Assessment Of The Buriganga River Using Qual2k. *Global Mainstream Journal of Innovation, Engineering & Emerging Technology*, *2*(03), 01-11. https://doi.org/10.62304/jieet.v2i03.64
[91]. Md Mahamudur Rahaman, S. (2022). Electrical And Mechanical Troubleshooting in Medical And Diagnostic Device Manufacturing: A Systematic Review Of Industry Safety And Performance Protocols. *American Journal of Scholarly Research and Innovation*, *1*(01), 295-318. https://doi.org/10.63125/d68y3590
[92]. Md Mahamudur Rahaman, S., & Rezwanul Ashraf, R. (2022). Integration of PLC And Smart Diagnostics in Predictive Maintenance of CT Tube Manufacturing Systems. *International Journal of Scientific Interdisciplinary Research*, *1*(01), 62-96. https://doi.org/10.63125/gspb0f75
[93]. Md Nazrul Islam, K. (2022). A Systematic Review of Legal Technology Adoption In Contract Management, Data Governance, And Compliance Monitoring. *American Journal of Interdisciplinary Studies*, *3*(01), 01-30. https://doi.org/10.63125/caangg06
[94]. Md Nur Hasan, M., Md Musfiqur, R., & Debashish, G. (2022). Strategic Decision-Making in Digital Retail Supply Chains: Harnessing AI-Driven Business Intelligence From Customer Data. *Review of Applied Science and Technology*, *1*(03), 01-31. https://doi.org/10.63125/6a7rpy62
[95]. Md Redwanul, I., & Md. Zafor, I. (2022). Impact of Predictive Data Modeling on Business Decision-Making: A Review Of Studies Across Retail, Finance, And Logistics. *American Journal of Advanced Technology and Engineering Solutions*, *2*(02), 33-62. https://doi.org/10.63125/8hfbkt70
[96]. Md Rezaul, K., & Md Mesbaul, H. (2022). Innovative Textile Recycling and Upcycling Technologies For Circular Fashion: Reducing Landfill Waste And Enhancing Environmental Sustainability. *American Journal of Interdisciplinary Studies*, *3*(03), 01-35. https://doi.org/10.63125/kkmerg16
[97]. Md Sultan, M., Proches Nolasco, M., & Md. Torikul, I. (2023). Multi-Material Additive Manufacturing For Integrated Electromechanical Systems. *American Journal of Interdisciplinary Studies*, *4*(04), 52-79. https://doi.org/10.63125/y2ybrx17
[98]. Md Sultan, M., Proches Nolasco, M., & Vicent Opiyo, N. (2025). A Comprehensive Analysis Of Non-Planar Toolpath Optimization In Multi-Axis 3D Printing: Evaluating The Efficiency Of Curved Layer Slicing Strategies. *Review of Applied Science and Technology*, *4*(02), 274-308. https://doi.org/10.63125/5fdxa722
[99]. Md Takbir Hossen, S., Ishtiaque, A., & Md Atiqur, R. (2023). AI-Based Smart Textile Wearables For Remote Health Surveillance And Critical Emergency Alerts: A Systematic Literature Review. *American Journal of Scholarly Research and Innovation*, *2*(02), 1-29. https://doi.org/10.63125/ceqapd08
[100]. Md Tawfiqul, I. (2023). A Quantitative Assessment Of Secure Neural Network Architectures For Fault Detection In Industrial Control Systems. *Review of Applied Science and Technology*, *2*(04), 01-24. https://doi.org/10.63125/3m7gbs97

[101]. Md. Sakib Hasan, H. (2022). Quantitative Risk Assessment of Rail Infrastructure Projects Using Monte Carlo Simulation And Fuzzy Logic. *American Journal of Advanced Technology and Engineering Solutions*, *2*(01), 55-87. https://doi.org/10.63125/h24n6z92

[102]. Md. Tarek, H. (2022). Graph Neural Network Models For Detecting Fraudulent Insurance Claims In Healthcare Systems. *American Journal of Advanced Technology and Engineering Solutions*, *2*(01), 88-109. https://doi.org/10.63125/r5vsmv21

[103]. Md.Kamrul, K., & Md Omar, F. (2022). Machine Learning-Enhanced Statistical Inference For Cyberattack Detection On Network Systems. *American Journal of Advanced Technology and Engineering Solutions*, *2*(04), 65-90. https://doi.org/10.63125/sw7jzx60

[104]. Md.Kamrul, K., & Md. Tarek, H. (2022). A Poisson Regression Approach to Modeling Traffic Accident Frequency in Urban Areas. *American Journal of Interdisciplinary Studies*, *3*(04), 117-156. https://doi.org/10.63125/wqh7pd07

[105]. Melaku, H. M. (2023). Context-based and adaptive cybersecurity risk management framework. *Risks*, *11*(6), 101.

[106]. Meyers, J., Fabian, B., & Brown, N. (2021). De novo molecular design and generative models. *Drug discovery today*, *26*(11), 2707-2715.

[107]. Min, E., Guo, X., Liu, Q., Zhang, G., Cui, J., & Long, J. (2018). A survey of clustering with deep learning: From the perspective of network architecture. *IEEE Access*, *6*, 39501-39514.

[108]. Mishra, A., Alzoubi, Y. I., Gill, A. Q., & Anwar, M. J. (2022). Cybersecurity enterprises policies: A comparative study. *Sensors*, *22*(2), 538.

[109]. Möller, D. P. (2023a). Cybersecurity in digital transformation. In *Guide to cybersecurity in digital transformation: Trends, methods, technologies, applications and best practices* (pp. 1-70). Springer.

[110]. Möller, D. P. (2023b). Ransomware attacks and scenarios: Cost factors and loss of reputation. In *Guide to cybersecurity in digital transformation: Trends, methods, Technologies, Applications and best practices* (pp. 273-303). Springer.

[111]. Mshragi, M., & Petri, I. (2025). Fast machine learning for building management systems. *Artificial intelligence review*, *58*(7), 211.

[112]. Mubashir, I., & Abdul, R. (2022). Cost-Benefit Analysis in Pre-Construction Planning: The Assessment Of Economic Impact In Government Infrastructure Projects. *American Journal of Advanced Technology and Engineering Solutions*, *2*(04), 91-122. https://doi.org/10.63125/kjwd5e33

[113]. Muñoz-La Rivera, F., Mora-Serrano, J., Valero, I., & Oñate, E. (2021). Methodological-technological framework for construction 4.0. *Archives of computational methods in engineering*, *28*(2), 689-711.

[114]. Nankya, M., Chataut, R., & Akl, R. (2023). Securing industrial control systems: Components, cyber threats, and machine learning-driven defense strategies. *Sensors*, *23*(21), 8840.

[115]. Nash, W., Drummond, T., & Birbilis, N. (2018). A review of deep learning in the study of materials degradation. *npj Materials Degradation*, *2*(1), 37.

[116]. Ngo, D., Park, H.-C., & Kang, B. (2025). Edge Intelligence: A Review of Deep Neural Network Inference in Resource-Limited Environments. *Electronics*, *14*(12), 2495.

[117]. Omar Muhammad, F., & Md.Kamrul, K. (2022). Blockchain-Enabled BI For HR And Payroll Systems: Securing Sensitive Workforce Data. *American Journal of Scholarly Research and Innovation*, *1*(02), 30-58. https://doi.org/10.63125/et4bhy15

[118]. Pollmeier, S., Bongiovanni, I., & Slapničar, S. (2023). Designing a financial quantification model for cyber risk: A case study in a bank. *Safety science*, *159*, 106022.

[119]. Pupentsova, S., & Livintsova, M. (2021). The enterprises risk management in the context of digital transformation. International Scientific Siberian Transport Forum,

[120]. Radanliev, P. (2024). The rise and fall of cryptocurrencies: defining the economic and social values of blockchain technologies, assessing the opportunities, and defining the financial and cybersecurity risks of the Metaverse. *Financial Innovation*, *10*(1), 1.

[121]. Radanliev, P., De Roure, D. C., Nicolescu, R., Huth, M., Montalvo, R. M., Cannady, S., & Burnap, P. (2018). Future developments in cyber risk assessment for the internet of things. *Computers in industry*, *102*, 14-22.

[122]. Rajawat, A. S., Goyal, S., Verma, C., & Singh, J. (2024). Advancing network security paradigms integrating quantum computing models for enhanced protections. In *Applied Data Science and Smart Systems* (pp. 517-528). CRC Press.

[123]. Rea-Guaman, A. M., Mejía, J., San Feliu, T., & Calvo-Manzano, J. A. (2020). AVARCIBER: a framework for assessing cybersecurity risks. *Cluster Computing*, *23*(3), 1827-1843.

[124]. Reduanul, H., & Mohammad Shoeb, A. (2022). Advancing AI in Marketing Through Cross Border Integration Ethical Considerations And Policy Implications. *American Journal of Scholarly Research and Innovation*, *1*(01), 351-379. https://doi.org/10.63125/d1xg3784

[125]. Rejeb, A., Rejeb, K., Zrelli, I., Kayikci, Y., & Hassoun, A. (2024). The research landscape of industry 5.0: a scientific mapping based on bibliometric and topic modeling techniques. *Flexible Services and Manufacturing Journal*, 1-48.

[126]. Rodriguez-Conde, I., Campos, C., & Fdez-Riverola, F. (2023). Horizontally distributed inference of deep neural networks for AI-enabled IoT. *Sensors*, *23*(4), 1911.

[127]. Sabuj Kumar, S., & Zobayer, E. (2022). Comparative Analysis of Petroleum Infrastructure Projects In South Asia And The Us Using Advanced Gas Turbine Engine Technologies For Cross Integration. *American Journal of Advanced Technology and Engineering Solutions*, *2*(04), 123-147. https://doi.org/10.63125/wr93s247
[128]. Sadia, T., & Shaiful, M. (2022). In Silico Evaluation of Phytochemicals From Mangifera Indica Against Type 2 Diabetes Targets: A Molecular Docking And Admet Study. *American Journal of Interdisciplinary Studies*, *3*(04), 91-116. https://doi.org/10.63125/anaf6b94
[129]. Safitra, M. F., Lubis, M., & Fakhrurroja, H. (2023). Counterattacking cyber threats: A framework for the future of cybersecurity. *Sustainability*, *15*(18), 13369.
[130]. Saleem, M. H., Potgieter, J., & Arif, K. M. (2021). Automation in agriculture by machine and deep learning techniques: A review of recent developments. *Precision Agriculture*, *22*(6), 2053-2091.
[131]. Sánchez-García, I. D., Mejía, J., & San Feliu Gilabert, T. (2022). Cybersecurity risk assessment: a systematic mapping review, proposal, and validation. *Applied Sciences*, *13*(1), 395.
[132]. Sanjai, V., Sanath Kumar, C., Maniruzzaman, B., & Farhana Zaman, R. (2023). Integrating Artificial Intelligence in Strategic Business Decision-Making: A Systematic Review Of Predictive Models. *International Journal of Scientific Interdisciplinary Research*, *4*(1), 01-26. https://doi.org/10.63125/s5skge53
[133]. Sanjai, V., Sanath Kumar, C., Sadia, Z., & Rony, S. (2025). AI And Quantum Computing For Carbon-Neutral Supply Chains: A Systematic Review Of Innovations. *American Journal of Interdisciplinary Studies*, *6*(1), 40-75. https://doi.org/10.63125/nrdx7d32
[134]. Sarker, I. H., Janicke, H., Maglaras, L., & Camtepe, S. (2023). Data-driven intelligence can revolutionize today's cybersecurity world: A position paper. International Conference on Advanced Research in Technologies, Information, Innovation and Sustainability,
[135]. Schöbel, S., Schmitt, A., Benner, D., Saqr, M., Janson, A., & Leimeister, J. M. (2024). Charting the evolution and future of conversational agents: A research agenda along five waves and new frontiers. *Information Systems Frontiers*, *26*(2), 729-754.
[136]. Schreiber, C., Abbad-Andaloussi, A., & Weber, B. (2023). On the cognitive effects of abstraction and fragmentation in modularized process models. International Conference on Business Process Management,
[137]. Sepasgozar, S. M. (2021). Differentiating digital twin from digital shadow: Elucidating a paradigm shift to expedite a smart, sustainable built environment. *Buildings*, *11*(4), 151.
[138]. Serpanos, D., & Wolf, M. (2018). Internet-of-Things (IoT) Systems. *Architectures, Algorithms, Methodologies*.
[139]. Shaikh, F. A., & Siponen, M. (2023). Information security risk assessments following cybersecurity breaches: The mediating role of top management attention to cybersecurity. *Computers & Security*, *124*, 102974.
[140]. Shandilya, S. K., Datta, A., Kartik, Y., & Nagar, A. (2024). Navigating the regulatory landscape. In *Digital Resilience: Navigating Disruption and Safeguarding Data Privacy* (pp. 127-240). Springer.
[141]. Sheratun Noor, J., & Momena, A. (2022). Assessment Of Data-Driven Vendor Performance Evaluation in Retail Supply Chains: Analyzing Metrics, Scorecards, And Contract Management Tools. *American Journal of Interdisciplinary Studies*, *3*(02), 36-61. https://doi.org/10.63125/0s7t1y90
[142]. Shrestha, A., & Mahmood, A. (2019). Review of deep learning algorithms and architectures. *IEEE Access*, *7*, 53040-53065.
[143]. Singh, S. P., Wang, L., Gupta, S., Goli, H., Padmanabhan, P., & Gulyás, B. (2020). 3D deep learning on medical images: a review. *Sensors*, *20*(18), 5097.
[144]. Sivanathan, A., Gharakheili, H. H., & Sivaraman, V. (2020). Managing IoT cyber-security using programmable telemetry and machine learning. *IEEE Transactions on Network and Service Management*, *17*(1), 60-74.
[145]. Sreenu, G., & Durai, S. (2019). Intelligent video surveillance: a review through deep learning techniques for crowd analysis. *Journal of big Data*, *6*(1), 1-27.
[146]. Srinivas, J., Das, A. K., & Kumar, N. (2019). Government regulations in cyber security: Framework, standards and recommendations. *Future generation computer systems*, *92*, 178-188.
[147]. Stetco, A., Dinmohammadi, F., Zhao, X., Robu, V., Flynn, D., Barnes, M., Keane, J., & Nenadic, G. (2019). Machine learning methods for wind turbine condition monitoring: A review. *Renewable energy*, *133*, 620-635.
[148]. Strupczewski, G. (2021). Defining cyber risk. *Safety science*, *135*, 105143.
[149]. Suganyadevi, S., Seethalakshmi, V., & Balasamy, K. (2022). A review on deep learning in medical image analysis. *International Journal of Multimedia Information Retrieval*, *11*(1), 19-38.
[150]. Taherdoost, H. (2022). Understanding cybersecurity frameworks and information security standards—a review and comprehensive overview. *Electronics*, *11*(14), 2181.
[151]. Tahmina Akter, R., Debashish, G., Md Soyeb, R., & Abdullah Al, M. (2023). A Systematic Review of AI-Enhanced Decision Support Tools in Information Systems: Strategic Applications In Service-Oriented Enterprises And Enterprise Planning. *Review of Applied Science and Technology*, *2*(01), 26-52. https://doi.org/10.63125/73djw422
[152]. Tariq, U., Ahmed, I., Bashir, A. K., & Shaukat, K. (2023). A critical cybersecurity analysis and future research directions for the internet of things: A comprehensive review. *Sensors*, *23*(8), 4117.
[153]. Taye, M. M. (2023a). Theoretical understanding of convolutional neural network: Concepts, architectures, applications, future directions. *Computation*, *11*(3), 52.

[154]. Taye, M. M. (2023b). Understanding of machine learning with deep learning: architectures, workflow, applications and future directions. *Computers*, *12*(5), 91.
[155]. Thayyib, P., Mamilla, R., Khan, M., Fatima, H., Asim, M., Anwar, I., Shamsudheen, M., & Khan, M. A. (2023). State-of-the-art of artificial intelligence and big data analytics reviews in five different domains: a bibliometric summary. *Sustainability*, *15*(5), 4026.
[156]. Tzavara, V., & Vassiliadis, S. (2024). Tracing the evolution of cyber resilience: a historical and conceptual review. *International Journal of Information Security*, *23*(3), 1695-1719.
[157]. Uddin, M. H., Ali, M. H., & Hassan, M. K. (2020). Cybersecurity hazards and financial system vulnerability: a synthesis of literature. *Risk Management*, *22*(4), 239-309.
[158]. Villalón-Fonseca, R. (2022). The nature of security: A conceptual framework for integral-comprehensive modeling of IT security and cybersecurity. *Computers & Security*, *120*, 102805.
[159]. Wang, M., Yang, N., Gunasinghe, D. H., & Weng, N. (2023). On the robustness of ML-based network intrusion detection systems: An adversarial and distribution shift perspective. *Computers*, *12*(10), 209.
[160]. Wang, Y., Han, Y., Wang, C., Song, S., Tian, Q., & Huang, G. (2024). Computation-efficient deep learning for computer vision: A survey. *Cybernetics and intelligence*.
[161]. Wang, Z., Zhang, H.-W., Dai, Y.-Q., Cui, K., Wang, H., Chee, P. W., & Wang, R.-F. (2025). Resource-Efficient Cotton Network: A Lightweight Deep Learning Framework for Cotton Disease and Pest Classification. *Plants*, *14*(13), 2082.
[162]. Wylde, V., Rawindaran, N., Lawrence, J., Balasubramanian, R., Prakash, E., Jayal, A., Khan, I., Hewage, C., & Platts, J. (2022). Cybersecurity, data privacy and blockchain: A review. *SN computer science*, *3*(2), 127.
[163]. Yadav, A., & Vishwakarma, D. K. (2020). Sentiment analysis using deep learning architectures: a review. *Artificial intelligence review*, *53*(6), 4335-4385.
[164]. Yamashita, R., Nishio, M., Do, R. K. G., & Togashi, K. (2018). Convolutional neural networks: an overview and application in radiology. *Insights into imaging*, *9*(4), 611-629.
[165]. Yao, G., Lei, T., & Zhong, J. (2019). A review of convolutional-neural-network-based action recognition. *Pattern Recognition Letters*, *118*, 14-22.
[166]. Zhang, L., Chen, Z., Laili, Y., Ren, L., Deen, M. J., Cai, W., Zhang, Y., Zeng, Y., & Gu, P. (2025). MBSE 2.0: Toward more integrated, comprehensive, and intelligent MBSE. *Systems*, *13*(7), 584.
[167]. Zhang, W., Li, H., Li, Y., Liu, H., Chen, Y., & Ding, X. (2021). Application of deep learning algorithms in geotechnical engineering: a short critical review. *Artificial intelligence review*, *54*(8), 5633-5673.
[168]. Zheng, Z., Wan, Y., Zhang, Y., Xiang, S., Peng, D., & Zhang, B. (2021). CLNet: Cross-layer convolutional neural network for change detection in optical remote sensing imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, *175*, 247-267.
[169]. Zhou, D.-X. (2020). Universality of deep convolutional neural networks. *Applied and computational harmonic analysis*, *48*(2), 787-794.