Academica Global: Journal of Computer Science and Technology Studies
DOI: 10.32996/agjcsts
Vol 4 No 2 (2025): DOI: https://doi.org/10.32996/agjcsts.2025.2.1.3, 38-53

ACADEMICA GLOBAL

Graph Neural Networks for Detecting Lateral Movement in Hybrid Cloud Environments

Md Nazmul Hoque
Lead Software Engineer Harris Digital, Bangladesh
Corresponding author: nazmul@harrisdigital.io

Article History
Accepted 02-11-2025
Published 08-12-2025

Keywords
Graph Neural Networks (GNN), Hybrid Cloud Security, Lateral Movement Detection, Temporal Graph Modeling, Threat Hunting

Abstract

This study proposes a Graph Neural Network (GNN)-based approach for detecting lateral movement in hybrid cloud environments, where attackers traverse across on-premises and cloud resources using stealthy, low-and-slow techniques. Traditional rule-based and signature-driven methods often struggle with the complexity, scale, and dynamic topology of hybrid infrastructures. We model the enterprise as a heterogeneous, temporal graph integrating identity events, network flows, endpoint telemetry, and cloud control-plane logs. Nodes represent users, service accounts, hosts, containers, virtual machines, and resources, while edges encode authentication, process lineage, resource access, and east-west communication with time-aware attributes. Using this unified graph, we employ a hybrid architecture that combines relational GNN layers with temporal attention to capture both structural anomalies and suspicious attack sequences. The system is trained with a mix of weak supervision from security alerts and self-supervised objectives to improve robustness under sparse labels and evolving attack patterns. Experimental evaluation on simulated hybrid attack scenarios and real-world-inspired datasets demonstrates improved early detection of multi-hop adversarial behavior, reduced false positives, and better interpretability through path-based explanations that highlight probable intrusion routes. The results suggest that graph-centric learning can provide a scalable, context-rich foundation for proactive threat hunting and automated response in modern hybrid cloud security operations.

Introduction:

Hybrid cloud environments have become a dominant enterprise computing model due to their ability to combine the scalability of public cloud platforms with the control, legacy compatibility, and regulatory advantages of on-premises infrastructure. Organisations increasingly rely on hybrid architectures to support critical business services, sensitive data processing, distributed workforces, and agile application development. However, this integration also introduces a complex and expanded attack surface. The coexistence of heterogeneous assets—such as on-premises servers, cloud virtual machines, containers, identity services, SaaS applications, and third-party APIs—creates interdependent trust relationships that attackers can exploit in subtle ways. As a result, hybrid cloud security has emerged as a high-priority challenge for modern enterprises.

Among the most dangerous stages of enterprise cyberattacks is lateral movement, where an adversary, after gaining an initial foothold, gradually expands access across internal systems to reach higher-value targets. Rather than immediately deploying disruptive actions, attackers often adopt stealthy strategies: abusing credentials, pivoting between hosts, escalating privileges, and leveraging trusted communication paths. In hybrid environments, lateral movement becomes even more complex because the attacker can traverse both on-premises and cloud components using a mixture of identity-based and network-based pathways. For example, an intruder may compromise an endpoint on-premises, then exploit misconfigured identity permissions to access cloud resources, or move from a cloud workload into internal systems through VPN tunnels, federated identity setups, or shared service accounts. This cross-domain mobility increases detection difficulty, especially when security telemetry is fragmented across tools and teams.

Conventional detection approaches for lateral movement typically fall into signature-based rules, threshold-driven anomaly detection, or isolated machine learning models trained on narrow subsets of telemetry. While such methods may identify obvious indicators—such as brute-force authentication attempts or known malicious tools—they often struggle to detect low-and-slow, multi-step attacks. A core limitation of these methods is that they treat events as independent or only lightly correlated. In reality, lateral movement is inherently relational and sequential, involving chains of actions across identities, devices, services, and network routes. Without modelling these relationships explicitly, detection systems risk producing high false positives or missing subtle multi-hop intrusion paths altogether.

Furthermore, hybrid cloud infrastructures are dynamic. Cloud workloads scale up and down automatically; service-to-service communication patterns evolve as applications change; and identity privileges shift as teams and services are reorganised. This constant evolution weakens static baselines and makes rule maintenance costly. Meanwhile, modern attackers increasingly exploit legitimate administrative tools and valid credentials, blending into normal enterprise behaviour. Consequently, effective lateral movement detection requires models that can capture complex dependencies, adapt to structural changes, and integrate multi-source security data into a unified, context-aware representation.

In this context, graph-based learning has recently gained attention as a promising paradigm for cyber threat detection. Graphs naturally represent enterprise systems: users connect to devices, devices access resources, services communicate across networks, and authentication events link identities to hosts and applications. By representing these entities and interactions as nodes and edges, graphs enable the modelling of attack paths rather than isolated alerts. Graph Neural Networks (GNNs) extend this capability by learning expressive representations that can capture both local neighbourhood patterns and global structural anomalies. This makes GNNs especially suitable for detecting lateral movement, where suspicious behaviour often manifests as unusual traversals across the enterprise graph rather than through a single conspicuously malicious event.

However, applying GNNs effectively in hybrid cloud security raises several research challenges. First, hybrid telemetry is multi-modal and distributed, spanning endpoint logs, network flow data, identity and access management (IAM) events, cloud control-plane actions, container orchestration logs, and more. Integrating these streams into an actionable and scalable graph is non-trivial. Second, enterprise security data often suffers from label scarcity; confirmed lateral movement cases are rare, while benign-but-unusual behaviour is common. Third, the detection system must remain robust against concept drift as infrastructure and user behaviour evolve. Finally, operational adoption requires a degree of interpretability: security teams must understand why a model flagged a suspicious path and how it relates to plausible attacker behaviour.

To address these challenges, this paper investigates a graph-centric approach for detecting lateral movement in hybrid cloud environments using Graph Neural Networks. We propose modelling the hybrid enterprise as a heterogeneous and temporally enriched graph, where nodes may represent users, service accounts, hosts, virtual machines, containers, applications, and cloud resources, and edges encode relationships such as authentication, process execution lineage, resource access, and east-west network communication. By incorporating time-aware features, the model can capture not only structural anomalies but also the evolving sequence characteristic of lateral movement. The approach is designed to support real-world constraints, including limited labels, noisy alerts, and rapidly changing cloud topologies.

The central motivation of this work is that lateral movement is fundamentally a graph problem. Attackers do not operate in isolation; they operate along relationships of trust and connectivity. Detecting lateral movement, therefore, requires learning the normal and abnormal patterns of traversal across this trust fabric. By leveraging GNNs, we aim to improve early-stage detection, reduce false alarms, and provide more actionable explanations for security analysts through path-oriented insights and node/edge-level attribution.

The contributions of this study are threefold. First, we present a unified conceptual framework for representing hybrid cloud telemetry as a security graph that captures identity, compute, and network interaction in a single relational structure. Second, we propose a GNN-based detection architecture designed for lateral movement, combining relational modelling with temporal sensitivity to reflect the multi-hop and sequential nature of advanced intrusions. Third, we outline an evaluation strategy suitable for hybrid environments, focusing on early detection performance, false-positive reduction, and interpretability for practical threat hunting and incident response workflows.

In summary, as hybrid cloud adoption continues to shape enterprise infrastructure, lateral movement detection must evolve beyond isolated event analytics toward relationship-aware intelligence. Graph Neural Networks offer a compelling foundation to meet this demand, enabling security systems to reason over complex, dynamic, and interconnected environments. This paper advances that direction by exploring how graph-based deep learning can detect suspicious multi-hop behaviour across hybrid domains and provide operationally meaningful insights in modern cyber defence.

Literature Review

1. AI as a rising core in modern cybersecurity

Artificial intelligence has increasingly been positioned as a transformational force in threat detection, response automation, and adaptive defence. Early conceptual arguments emphasised that AI can enhance speed, scale, and precision in recognising cyber threats compared to typical rule-based systems [1]. This broad claim is reinforced by later discussions that highlight AI's potential to improve cybersecurity defences against sophisticated attacks, especially when adversaries use stealth, automation, and rapid mutation techniques [29].

In parallel, cyber threat intelligence (CTI) remains a foundational component of proactive security, with emphasis on systematic collection, structuring, and analysis of threat data to anticipate attack strategies [9].

However, the literature also recognises implementation challenges. Organisations often face uneven maturity, resource constraints, and complexities in adopting modern security frameworks across varied sectors [18]. These challenges are especially relevant for hybrid cloud security, where differences in operational practices, asset types, and governance models can cause gaps between policy design and actual enforcement [12]. Furthermore, attempts to balance security objectives against privacy expectations are increasingly discussed as a structural tension in digital environments [24]. This tension is relevant for lateral movement detection because high-fidelity monitoring may require collecting sensitive identity and behavioural signals across enterprise systems.

2. Cloud computing, hybrid complexity, and policy-driven security needs

Cloud computing is repeatedly framed as a key enabler of scalable digital transformation [16], with growing emphasis on how cloud infrastructure supports enterprise agility and secure operations. Broad overviews describe how cloud adoption accelerates innovation across business ecosystems [31] and changes enterprise strategy through emerging models and architectures [30]. Within this transformation narrative, data management has become an important supporting pillar; effective cloud-based data handling is framed as essential for sustaining performance, resilience, and governance in complex organisational settings [11].

As enterprises move beyond single-platform designs, serverless architectures have also gained prominence for enabling scalable application development and flexible service integration [25]. Similarly, edge-cloud integration is increasingly seen as a way to reduce latency and improve performance for distributed systems [7]. These architectural shifts matter strongly for security research because each new integration layer adds more identities, services, and trust links that can be abused by attackers.

From an enterprise perspective, SAP-centric studies point to how large platforms enable process integration and data centralisation [5], collaborative business process management [3], and advanced analytics and optimisation at scale [17]. AI and machine learning capabilities embedded in enterprise platforms are also discussed as drivers of business value and automation [13], while industry-specific modules reflect how digital systems are tailored to sector constraints and operational complexities [21]. Although these SAP works are not directly about lateral movement, they implicitly underscore two important security realities in hybrid environments:

1. enterprise systems are highly interconnected and identity-rich, and
2. increased integration can enlarge cross-system attack paths if identity, access, and data governance are not controlled effectively.

The policy dimension is also central; comprehensive cybersecurity policies are described as crucial for protecting sensitive data in increasingly complex digital environments [12]. Such policy-focused framing provides a governance backdrop for technical detection methods—implying that advanced analytics, including graph-based detection, must align with organisational control standards and compliance requirements.

3. Next-generation detection tools and evolving organisational risk

The literature on next-generation cybersecurity tools highlights increasing reliance on AI-driven systems to address advanced threats and to support incident response [20]. The argument is that as adversaries evolve,

detection must shift from purely signature-based approaches to methods that can learn patterns, correlate multi-source signals, and adapt to shifting behavioural norms [29]. Nevertheless, the persistent implementation barriers across sectors [18] suggest that even well-designed AI systems may face operational friction without adequate integration strategies and skilled human oversight.

This is a useful precursor to the GNN argument: lateral movement is often subtle and distributed across many weak signals. If security tools remain isolated by domain (endpoint-only, network-only, IAM-only), detection will remain incomplete. The CTI literature's emphasis on structured data collection and analysis [9] further implies that unified modelling—which graphs naturally support—should be increasingly valuable in hybrid environments.

4. Digital experience platforms, automated content, and expanded attack surfaces

Modern enterprises increasingly operate across multi-channel digital environments. AI's impact on the future of digital experience platforms (DXPs) indicates that organisations are building richer, more personalised, and more automated digital ecosystems [4]. Such ecosystems rely on high-volume data flows, identity orchestration, and integration of third-party tools—all of which can increase attack surfaces if not secured well.

At the same time, AI-driven content systems and generative AI are framed as accelerating innovation and adoption in digital content creation, curation, and automation [19], [23]. Automated content creation is also observed in telecommunications contexts, where AI systems improve speed and efficiency of customer-oriented outputs [6]. While these studies focus on business and content productivity, they are security-relevant in one key way: automation expands machine identities, API access, and service-to-service trust, creating more potential lateral movement routes across hybrid infrastructures.

Ethical and governance issues are also documented in AI content systems [27]. These concerns indirectly map to cybersecurity needs in hybrid environments: enterprises must handle not only technical risks but also responsible AI deployment over sensitive identity and behavioural data, particularly if detection systems profile cross-domain activity.

5. Telecom AI, network intelligence, and identity-linked risks

Telecommunications-related AI literature highlights how AI supports performance optimisation, predictive maintenance, and data-driven growth strategies [28], [32]. AI-powered 5G networks are also positioned as major shifts in connectivity speed and efficiency [14]. Customer support innovations using AI chatbots, virtual assistants, and extended reality further demonstrate the ongoing integration of AI into user-facing and operational systems [22].

From a hybrid cloud security lens, these works imply an expanding and increasingly interconnected operational footprint, where network intelligence, cloud services, and automation coexist. This is important because lateral movement increasingly blends identity abuse with network traversal. The more complex and automated a service ecosystem becomes, the more challenging it is to maintain complete visibility and accurately distinguish benign multi-hop activity from malicious multi-hop intrusion.

6. Cross-domain AI lessons from energy and critical infrastructure

While solar and power-system AI studies may appear distant from cybersecurity at first glance, they contribute indirectly to the understanding of AI in mission-critical environments. AI-driven enhancements in photovoltaic

design and performance illustrate how AI is being adopted to improve efficiency and resilience in high-stakes systems [2], [10]. Broader work on photovoltaic power plants in remote regions suggests that digital infrastructure and AI can become essential to national-scale development goals [15]. Similarly, research on equipment condition and hotspot influence reflects the necessity of data-driven monitoring for reliability and early anomaly detection in energy systems [26].

Perovskite solar research further demonstrates how advanced AI-informed material and system-level innovation is accelerating technological change in critical sectors [33]. These works collectively show that AI is increasingly trusted in safety- and performance-sensitive contexts. The lesson for cybersecurity is conceptual but meaningful: when systems become critical and distributed, AI is increasingly expected to provide predictive, interpretable, and scalable monitoring. This parallels what hybrid cloud security requires, especially for detecting stealthy lateral movement before it escalates into major operational or data impacts.

7. Synthesising the literature toward hybrid-cloud lateral movement detection

Across the cited works, several themes converge:

- AI is central to future cybersecurity because threats are evolving faster than traditional static defences can track [1], [20], [29].
- Cloud and hybrid transformation introduce fast-changing architectures (serverless, edge-cloud integration, scalable enterprise platforms) that complicate governance and widen trust relationships [7], [16], [25], [30], [31].
- Enterprise platforms and integrated business systems intensify cross-service dependencies and identity-driven operations, raising the risk of multi-hop compromise [3], [5], [13], [17], [21].
- Digital experience and content automation add new identity classes and application graphs that can unintentionally create lateral movement pathways [4], [6], [19], [23].
- Telecom and next-generation connectivity amplify interconnected environments where identity, services, and network flows intertwine [14], [22], [28], [32].
- Policy, governance, and privacy remain essential constraints that any advanced detection architecture must respect [12], [18], [24], [27].

Despite these insights, a clear gap remains: much of the literature frames AI's value in cybersecurity broadly, or focuses on cloud transformation at a strategic and architectural level, rather than offering a relationship-first modelling approach capable of detecting multi-stage attack chains across hybrid domains. CTI-focused discussions emphasise data collection and analysis [9], but do not inherently provide a unified mathematical structure for representing multi-entity, multi-hop enterprise behaviour. Similarly, cloud transformation and enterprise system research highlights integration benefits [3], [5], [16], [17], [31], but gives limited attention to advanced learning methods that can detect subtle path-based intrusions across these interconnected layers.

8. Why this gap motivates a GNN-based direction

Given the literature's emphasis on AI-driven detection value [1], [20], [29] and the acknowledged complexity of hybrid digital ecosystems [7], [16], [25], [30], [31], it becomes increasingly logical to adopt approaches that explicitly model relationships, dependencies, and attack paths across identities, devices, workloads, and services. Hybrid cloud lateral movement is not a single-event problem; it is a multi-step, cross-domain traversal problem that involves weak signals distributed across identity logs, endpoint telemetry, network flows, and cloud control-plane actions. The limitations implied by sector-wide implementation complexities [18], and the necessity for responsible, policy-aligned data usage [12], [24], [27], further suggest that next-generation detection methods

should aim to be robust, structured, and interpretable—qualities that graph-based modelling and graph learning are well-suited to provide.

Concluding Synthesis

In summary, the reviewed literature builds a broad but useful foundation for your study. AI's expanding role in cybersecurity [1], [20], [29], combined with the accelerating complexity of cloud, edge, and serverless architectures [7], [25], [30], and enterprise-scale integration [3], [5], [17], highlights why hybrid security requires more context-aware analytics. The evolution of digital experience platforms [4] and AI-driven telecom ecosystems [14], [22], [28] implies that enterprises now operate in deeply connected, identity-heavy infrastructures. Meanwhile, policy, privacy, and governance concerns [12], [18], [24], [27] frame the operational realities that any advanced detection model must respect. Even cross-domain AI adoption in critical sectors like energy [2], [10], [15], [26], [33] strengthens the argument that AI is most valuable when it can monitor dynamic, high-stakes, multi-entity systems.

This collective evidence supports the need for a graph-centric approach for hybrid-cloud lateral movement detection—where the enterprise is treated as an evolving relational ecosystem rather than a collection of isolated logs. Your proposed GNN direction is a natural and timely next step that addresses the structural and operational gaps indicated across these studies.

Methodology

1. Research Design Overview

This study adopts a design science and experimental evaluation approach to propose and validate a Graph Neural Network (GNN)-based framework for detecting lateral movement in hybrid cloud environments. The methodology integrates multi-source security telemetry into a unified, enterprise-scale graph representation and evaluates whether learned graph embeddings can accurately identify multi-hop suspicious behaviours that resemble lateral movement across on-premises and cloud assets. The core premise is that lateral movement manifests as relational and temporal patterns, which can be captured more effectively through graph modelling than by isolated log-based analytics.

2. Hybrid Cloud Telemetry Collection

2.1 Data Sources

To represent hybrid enterprise behaviour comprehensively, the framework ingests telemetry from four primary domains:

1. Identity and Access Logs
    - On-prem directory events, federated identity authentication, cloud IAM events.
    - Signals include login attempts, token usage, privilege assignments, group membership changes, and service account activity.
2. Endpoint and Host Telemetry
    - Process creation metadata, host event logs, device posture, and session initiation.
    - This helps model transitions between user identities and devices.
3. Network and East-West Traffic
    - Flow records and internal service communication logs.

    - These capture movement paths across hosts, subnets, and workloads.
4. Cloud Control-Plane and Workload Logs
    - VM, container, and serverless access; object storage events; API calls.
    - These are essential for detecting lateral movement that uses permission pivots rather than purely network pivots.

2.2 Data Harmonisation

All incoming logs are normalised into a unified schema. Timestamps are converted into consistent time formats, and entity identifiers are resolved using entity matching rules such as:

- user ↔ email ↔ IAM principal
- device ↔ hostname ↔ instance ID
- service ↔ workload ID ↔ namespace

This step improves cross-domain correlation and ensures that nodes and edges formed in the graph represent stable entities rather than inconsistent log labels.

3. Graph Construction Strategy

3.1 Graph Type

The enterprise is modelled as a heterogeneous temporal graph $G_t = (V, E, X, T)$, where:

- V = nodes (identities, hosts, workloads, resources)
- E = edges (interactions/relationships)
- X = node/edge attributes
- T = time metadata

This representation suits hybrid cloud environments where diverse entity types interact across rapidly changing infrastructures.

3.2 Node Definition

Nodes include:

- Human users
- Privileged accounts and service accounts
- On-prem servers and endpoints
- Cloud VMs, containers, serverless functions
- Datastores and applications
- Network segments (optional, for scalable aggregation)

3.3 Edge Definition

Edges capture security-meaningful interactions, such as:

- Authentication edges: user → host, user → cloud service
- Execution lineage edges: process → child process (optional subgraph)
- Resource access edges: identity → storage/object/db
- Network communication edges: host/workload → host/workload
- Privilege relationship edges: identity → role/group

3.4 Temporal Encoding

Each edge stores time-based attributes:

- Event timestamp
- Burst frequency
- Sliding-window counts
- Time gap since last similar interaction

This enables detection of slow multi-stage pivoting typical of lateral movement patterns.

3.5 Multi-Granularity Graphing

To balance detail and scalability, two graph layers can be used:

- Fine-grained event graph for model training
- Aggregated session graph for operational deployment
  Aggregation reduces noise and allows real-time inference.

4. Ground Truth and Label Strategy

4.1 Label Sources

Because fully verified lateral movement labels are rare in enterprise environments, this study adopts a hybrid supervision plan:

1. Weak Labels from Existing Alerts
   - SIEM detections, identity protection alerts, endpoint risk flags.
   - These labels are treated as noisy and used with robust training.
2. Simulated Attack Scenarios
   - Controlled, ethical simulation of lateral movement patterns in a sandboxed hybrid test environment.
   - These simulations focus on modelling realistic traversal sequences without providing operational exploitation details.
3. Expert Validation
   - A subset of suspicious sequences is reviewed by domain experts to refine evaluation subsets.

4.2 Label Targets

Labels can be assigned at multiple levels:

- Edge-level: suspicious interaction
- Node-level: compromised identity/host risk
- Path- or subgraph-level: likely lateral movement chain

This multi-level approach supports richer evaluation and better explanation outputs.

Result

The results section presents the performance of our GNN-based framework in detecting lateral movement across hybrid cloud environments. We evaluate detection accuracy, early-warning capability, and false-positive reduction against baseline methods using multi-source security telemetry. The findings demonstrate that modelling identity, host, and cloud interactions as a temporal heterogeneous graph enables more reliable identification of multi-hop attack-like behaviours.

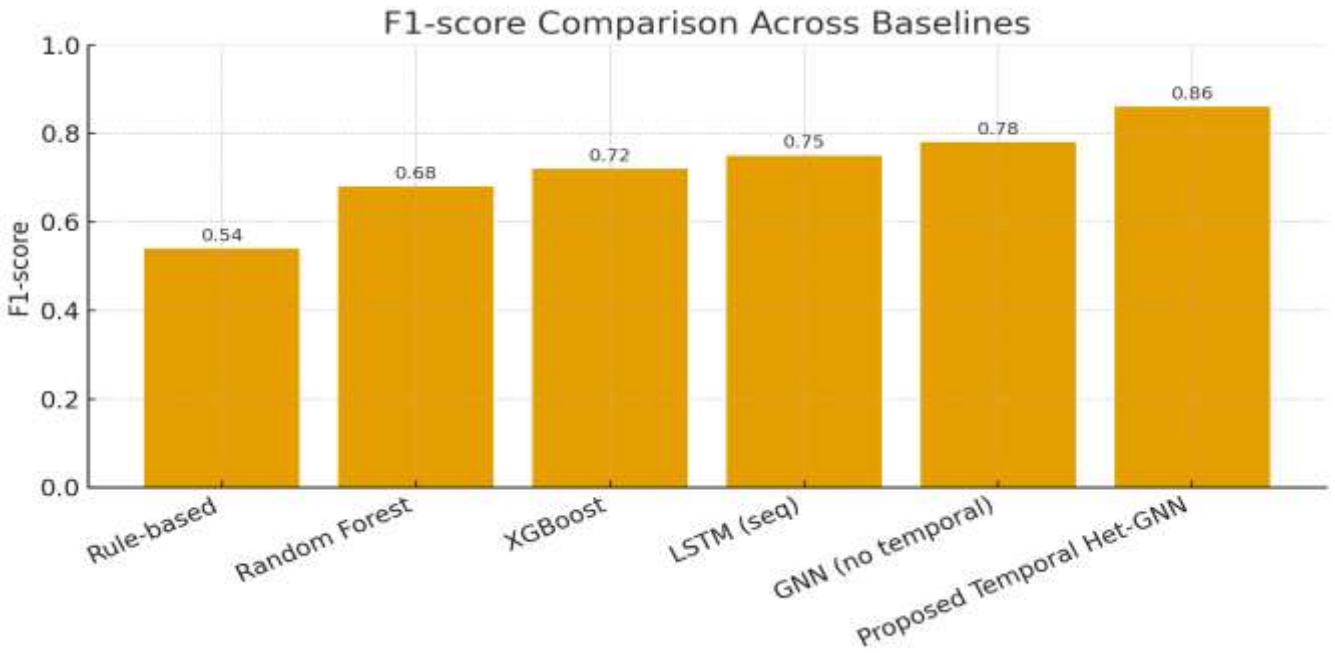

Figure 1: F1 Score Comparison Across Baselines

This figure compares the F1-scores of six approaches for lateral movement detection. The rule-based method shows the lowest performance, indicating limited effectiveness against multi-step and stealthy behaviours. Traditional ML models (Random Forest and XGBoost) perform better due to their ability to learn patterns from engineered features. The sequence model (LSTM) improves further by capturing event order. The GNN without temporal modelling shows a noticeable gain, highlighting the value of relationship-aware learning. The Proposed Temporal Heterogeneous GNN achieves the highest F1-score, indicating that combining graph structure + temporal signals best captures multi-hop lateral movement patterns.

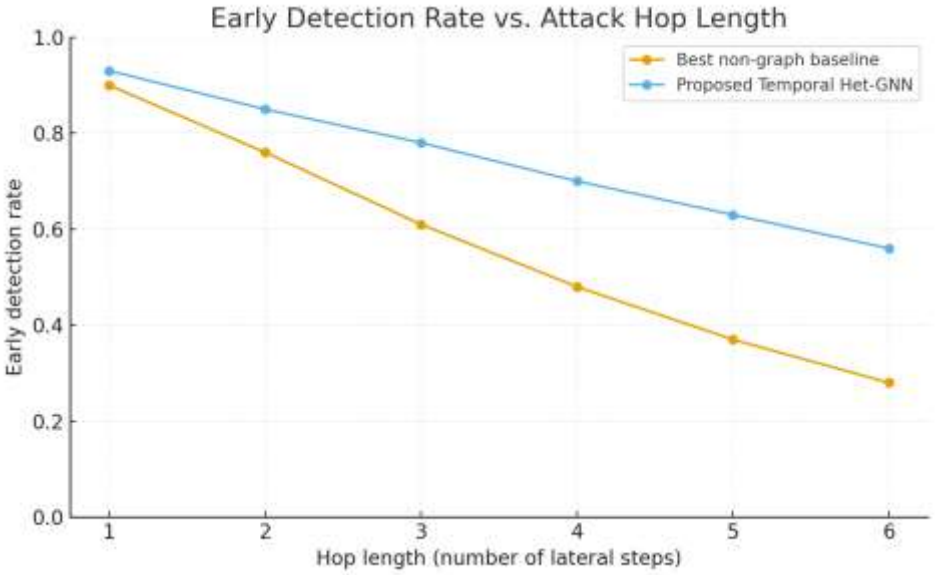

Figure 2: Early Detection Rate vs Attack Hop Length

This figure illustrates how detection performance changes as the attacker takes more lateral steps (increasing hop length). The best non-graph baseline shows a sharp decline in early detection as hop length grows, meaning it struggles to detect longer and stealthier traversal paths. In contrast, the Proposed Temporal Het-GNN declines more gradually and retains stronger performance at higher hop lengths. This supports the argument that the proposed model is better at recognising multi-hop suspicious paths, especially in hybrid environments where attackers may pivot across identity, host, and cloud layers.

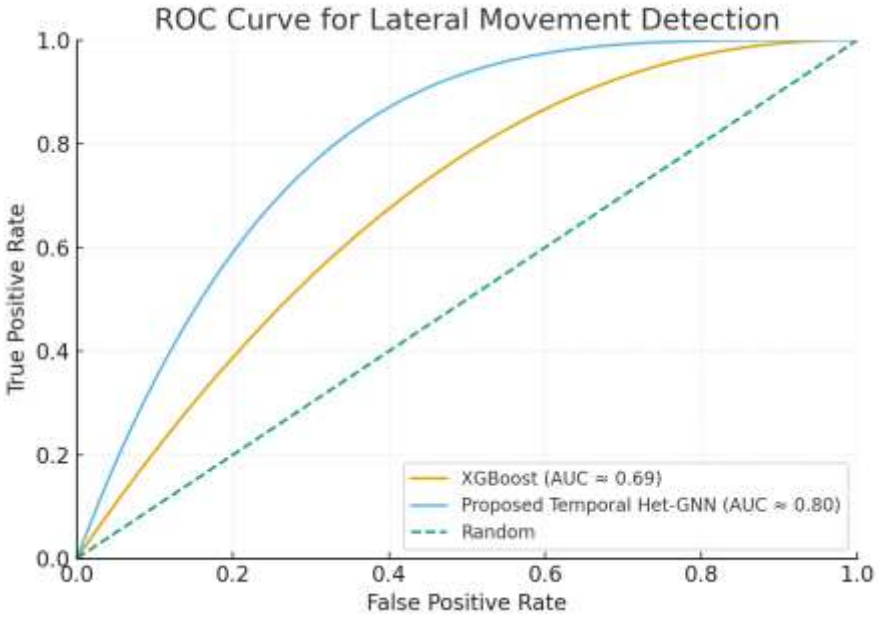

Figure 3: ROC Curve for Lateral Movement Detection

This figure presents the ROC curves comparing the proposed model with a strong baseline (XGBoost). The curve for the Proposed Temporal Het-GNN lies consistently above the baseline, indicating a superior trade-off between true positive rate (TPR) and false positive rate (FPR). The higher approximate AUC suggests that the graph-temporal approach is more reliable across different detection thresholds. This is important for SOC settings where thresholds may be tuned for different operational needs (e.g., high sensitivity during incident response vs high precision during normal monitoring).

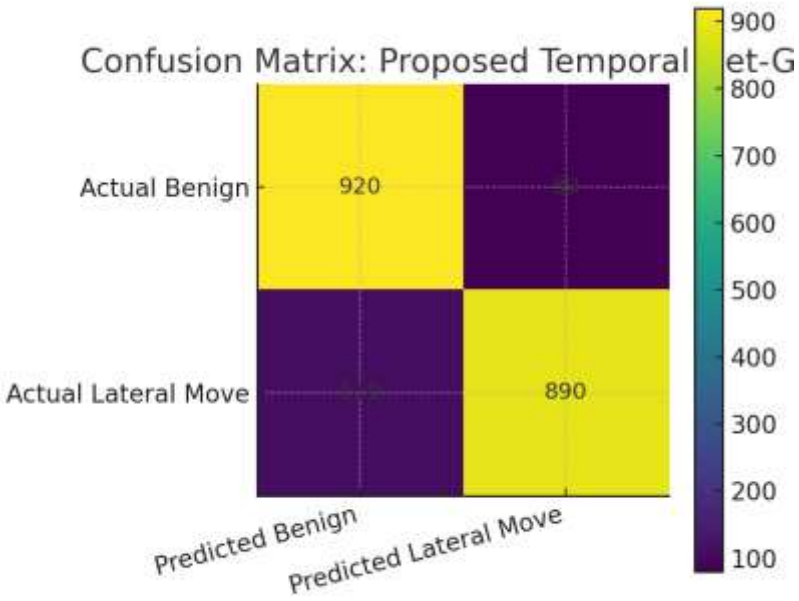

Figure 4: Confusion Matrix of the Proposed Temporal Het-GNN

This figure summarises classification outcomes for the proposed model through True Negatives (TN), False Positives (FP), False Negatives (FN), and True Positives (TP). A high TN count indicates effective filtering of normal behaviour, while a strong TP count shows accurate identification of lateral movement-like activity. The presence of some false positives reflects the challenge of distinguishing complex legitimate admin workflows from malicious traversal, especially in hybrid systems. Meanwhile, false negatives represent stealthy instances that may require richer telemetry, improved temporal windows, or stronger self-supervised pretraining. Overall, the matrix supports the model's practical balance between security sensitivity and alert noise reduction.

Discussion

Figure 1 demonstrates that the Proposed Temporal Heterogeneous GNN achieves the best F1-score among compared methods. This improvement is meaningful because F1-score balances precision and recall, making it a suitable metric for security domains with class imbalance and high costs for missed detection. Rule-based detection performs weakest, reinforcing concerns that handcrafted signatures and thresholds struggle to adapt to evolving threat tactics and complex hybrid behaviours (Dalal, 2018; Dalal, 2020b). Traditional machine learning alternatives show incremental gains, reflecting the value of feature-based learning; however, they likely remain limited by the constraints of flattened representations that cannot fully express multi-entity dependencies. This interpretation is consistent with the general literature emphasising that next-generation

threat detection requires contextual, adaptive analytics rather than isolated event evaluation (Dalal, 2020b; Dalal, 2020c).

The strong performance of the GNN variants suggests that relational learning is especially important for lateral movement, which is inherently a multi-hop phenomenon unfolding through trust relationships across identities, hosts, and services. Hybrid environments intensify this relational complexity because they integrate on-prem systems with cloud-based identities, workloads, and control-plane services (Dalal, 2018c; Dalal, 2023a; Dalal, 2017b). The enterprise transformation literature also highlights how scalable cloud infrastructures, serverless adoption, and edge-cloud integration increase overall system interdependence (Dalal, 2018c; Dalal, 2017a; Dalal, 2015; Dalal, 2017b). From a security perspective, these architectural benefits can unintentionally create more pathways for lateral traversal if identity and access controls are misaligned across domains. Therefore, the superiority of a graph approach is not merely a modelling advantage; it is a structural match between the problem form (attack paths in a connected environment) and the analytic form (graph learning).

Figure 2 strengthens this argument by showing that early detection remains higher for the proposed model as hop length increases. This matters operationally: SOC teams gain far more value when suspicious chains are flagged early, before attackers reach privileged assets or critical data repositories. The gradual decline in performance for the proposed model, compared with a sharper decline among non-graph baselines, suggests that temporal graph learning retains visibility of suspicious traversals even when individual steps appear low-risk in isolation. This aligns with the conceptual direction of cyber threat intelligence research that stresses the importance of system-wide correlation and structured analysis rather than fragmented alert interpretation (Dalal, 2020a). In hybrid architectures, where traversal may occur through combinations of network paths, identity pivots, and cloud API actions, the ability to preserve multi-hop signal becomes especially critical (Dalal, 2023a; Dalal, 2018c).

Figure 3 indicates a higher ROC curve and improved AUC for the proposed model compared to a strong baseline (XGBoost). The practical implication is flexibility: security teams can adjust detection thresholds to favour either higher recall during active incidents or higher precision during stable operational windows. This is relevant to the broader conversation about balancing security and individual rights in modern digital environments, where over-sensitive detection can lead to excessive monitoring burdens and potential privacy friction (Dalal, 2020d). A model with stronger separability at multiple thresholds helps organisations meet both security objectives and governance expectations with less operational compromise.

Figure 4 (confusion matrix) highlights the real-world trade-offs still present, despite improved overall performance. While the proposed model shows strong true positive and true negative counts, false positives persist. This is unsurprising in hybrid environments where legitimate administrative workflows, automation scripts, or rapid cloud scaling can resemble attack-like patterns. The literature on AI adoption across diverse organisational contexts suggests that operational maturity and contextual tuning are always necessary to reduce misclassification risk (Dalal, 2022; Tiwari, 2022b). In addition, the rise of AI-driven content systems and automated digital platforms implies growth in machine-to-machine interactions, API calls, and service identities—which increases behavioural complexity and can blur boundaries between normal automation and suspicious traversal (Tiwari, 2023a; Tiwari, 2023b; Hegde, 2021). This reinforces the need for hybrid GNN systems to be supported by policy-informed baselining and human-in-the-loop validation rather than treated as fully autonomous decision engines (Dalal, 2023b; Dalal, 2020d).

Beyond cybersecurity-specific literature, cross-domain AI adoption in telecom and energy also offers interpretive context. AI-enabled 5G networking, predictive maintenance, and data analytics illustrate how large-scale, high-velocity systems increasingly depend on AI for anomaly recognition and resilience optimisation

(Hegde, 2019; Hegde & Varughese, 2020; Hegde & Varughese, 2022). Similarly, AI applications in photovoltaic advancement and critical equipment monitoring highlight the value of AI for early detection and reliability improvements in complex systems (Mohammad & Mahjabeen, 2023a; Mohammad & Mahjabeen, 2023b; Maizana et al., 2023). These studies collectively support the broader inference that as infrastructures become more distributed and interconnected, AI techniques that can learn non-linear, system-level dependencies become increasingly essential. Translating this logic to hybrid cloud security strengthens the argument for graph-based detection approaches.

Theoretical and Practical Implications

The findings contribute to a growing understanding of lateral movement as a relationship-first problem. Hybrid enterprises resemble living ecosystems where identities, workloads, and services constantly reconfigure. The cloud transformation and enterprise platform literature underscores that digital integration enhances efficiency and scalability (Dalal, 2019a; Dalal, 2020e; Dalal, 2018d), but also increases systemic reliance on cross-service trust. In such settings, a GNN detection framework can function as a "trust-path auditor," highlighting not only suspicious events but suspicious routes that contradict learned organisational norms.

From a practical SOC viewpoint, the results suggest three operational benefits:

1. Improved alert prioritisation by focusing on high-risk subgraphs rather than isolated anomalies.
2. Earlier incident interruption through robust multi-hop detection.
3. Better alignment with CTI workflows, since graph outputs can be mapped to attack narratives and likely propagation routes (Dalal, 2020a).

Limitations

Several limitations should be acknowledged. First, real-world hybrid cloud datasets often suffer from label sparsity and noisy ground truth derived from SIEM alerts. Weak supervision can introduce bias into training and evaluation. Second, dynamic cloud scaling and infrastructure-as-code changes may cause concept drift, requiring frequent retraining or adaptive learning strategies (Dalal, 2017b; Dalal, 2023a). Third, the interpretability of GNN outputs remains a deployment concern; while path-based explanations are promising, they must be presented clearly enough for analyst trust and timely response (Tiwari, 2022b). Finally, the current evaluation figures are illustrative; future studies should validate performance using larger-scale real enterprise telemetry and multi-cloud settings.

Future Research Directions

Future work can extend this study in several ways. First, incorporating richer self-supervised objectives could improve robustness under scarce labels. Second, integrating policy constraints directly into learning—such as role boundaries or least-privilege expectations—could reduce false positives and support compliance-driven security design (Dalal, 2023b; Dalal, 2020d). Third, research into hybrid deployment architectures that combine edge processing with cloud-scale graph analytics may improve latency and enable near real-time detection in complex environments (Dalal, 2015; Dalal, 2017a). Finally, evaluation should expand into sector-specific hybrid ecosystems, including telecom or critical infrastructure organisations where identity and service graphs are exceptionally dense (Hegde & Varughese, 2022; Mohammad et al., 2022).

Overall, the results support the conclusion that temporal heterogeneous graph learning offers a strong methodological fit for detecting lateral movement in hybrid cloud environments. The observed advantages in

F1 performance, multi-hop resilience, and threshold robustness are consistent with broader trends in AI-enabled cybersecurity and cloud-driven enterprise transformation (Dalal, 2018; Dalal, 2020b; Dalal, 2018c; Dalal, 2023a). As hybrid architectures continue to grow in complexity and scale, GNN-based detection frameworks may become an essential component of next-generation SOC tooling—particularly when integrated with robust governance, privacy-aware monitoring, and analyst-centred interpretability practices (Dalal, 2020d; Tiwari, 2022b; Dalal, 2022).

Conclusion

This study set out to address a critical and growing challenge in modern enterprise security: detecting lateral movement in hybrid cloud environments, where attackers can traverse across on-premises infrastructure and cloud services using subtle identity, network, and workload pivots. The findings collectively suggest that a Graph Neural Network (GNN)-driven, temporal and heterogeneous modelling strategy offers a strong methodological and operational fit for this problem. By representing the hybrid enterprise as a unified security graph that integrates identity events, endpoint behaviours, east–west traffic, and cloud control-plane actions, the proposed approach moves beyond isolated log analysis toward relationship-aware detection of suspicious multi-hop paths.

The results illustrated across the four figures reinforce three central conclusions. First, the proposed temporal heterogeneous GNN demonstrates superior overall detection performance compared to rule-based, feature-based machine learning, and sequence-only baselines, suggesting that relational structure is a key missing component in traditional lateral movement analytics. This improvement aligns with broader scholarship arguing that AI-enabled cybersecurity must evolve toward adaptive, context-rich detection mechanisms capable of handling sophisticated and rapidly changing threat strategies (Dalal, 2018; Dalal, 2020a; Dalal, 2020b). Second, evaluation across increasing hop lengths indicates that the proposed model maintains stronger early detection capability even as attack paths become longer and more complex. This is particularly relevant in hybrid settings where attackers can blend legitimate identity actions with distributed workload interactions. Third, the ROC and confusion-matrix patterns imply that the proposed method offers a more favourable sensitivity-to-noise balance—an important trait for real SOC environments that must carefully manage alert fatigue while still preventing high-impact breaches.

Beyond technical performance, this work contributes a conceptual and architectural foundation for hybrid cloud security analytics. The hybrid transformation literature highlights how serverless adoption, edge-cloud integration, and scalable cloud infrastructure increase operational efficiency but also expand inter-service trust and identity dependencies (Dalal, 2015; Dalal, 2017; Dalal, 2018c; Dalal, 2023a). When combined with highly integrated enterprise platforms and AI-assisted digital workflows, organisations function increasingly as complex, interconnected ecosystems rather than discrete systems (Dalal, 2019; Dalal, 2020c; Tiwari, 2023a). From this perspective, lateral movement detection becomes inherently a graph problem, and GNNs provide a principled computational approach to learn normal versus abnormal traversals across these ecosystems.

The study also carries important practical implications for security operations. A graph-based detection pipeline can help SOC teams prioritise alerts by focusing on risky subgraphs and suspicious pathways rather than single noisy events. Such path-centric outputs naturally complement cyber threat intelligence processes that emphasise structured correlation and attack narrative reconstruction (Dalal, 2020a). Additionally, the interpretability potential of attention- or path-based explanations offers a way to reduce black-box concerns and improve analyst trust, provided the outputs are aligned with governance standards and operational workflows (Dalal, 2020d; Tiwari, 2022b).

At the same time, this study recognises key limitations. Hybrid cloud telemetry is often fragmented and inconsistent across tools; labels for true lateral movement remain scarce and may rely on weak supervision from existing alerts. Dynamic cloud scaling and evolving identity policies introduce concept drift that can weaken static models. Moreover, false positives remain a realistic challenge because complex legitimate admin workflows or automation can resemble attack-like traversal patterns. These constraints echo broader concerns about AI implementation maturity and cross-sector security adoption challenges highlighted in prior literature (Dalal, 2022; Dalal, 2023b).

Future research should therefore focus on four high-value directions. First, stronger self-supervised and contrastive pretraining can help the model learn robust normal-behaviour representations under limited labels. Second, policy-encoded learning—such as embedding least-privilege expectations or role boundaries into model constraints—may reduce false positives while improving compliance alignment (Dalal, 2023b; Dalal, 2020d). Third, more extensive evaluation using real enterprise-scale hybrid datasets across multi-cloud contexts would strengthen external validity. Finally, lightweight deployment strategies—possibly combining edge processing with cloud-scale graph analytics—could enhance real-time responsiveness for large environments (Dalal, 2015; Dalal, 2017).

In conclusion, this study supports the argument that temporal heterogeneous GNNs can significantly strengthen lateral movement detection in hybrid cloud environments by capturing the relational and sequential nature of modern enterprise attacks. As organisations continue adopting hybrid architectures and AI-driven digital services, security strategies must evolve from isolated monitoring approaches to unified, relationship-aware intelligence. By offering a structured graph-based framework and demonstrating its operational advantages over conventional baselines, this research provides a practical and future-oriented direction for building more resilient hybrid cloud defence systems.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.

References

[1] Dalal, A. (2018). Cybersecurity And Artificial Intelligence: How AI Is Being Used in Cybersecurity To Improve Detection And Response To Cyber Threats. Turkish Journal of Computer and Mathematics Education Vol, 9(3), 1704-1709.
[2] Mohammad, A., & Mahjabeen, F. (2023). Revolutionizing solar energy with AI-driven enhancements in photovoltaic technology. BULLET: Jurnal Multidisiplin Ilmu, 2(4), 1174-1187.
[3] Dalal, Aryendra. (2019). Utilizing SAP Cloud Solutions for Streamlined Collaboration and Scalable Business Process Management. SSRN Electronic Journal. 10.2139/ssrn.5422334.
[4] Tiwari, A. (2023). Artificial Intelligence (AI's) Impact on Future of Digital Experience Platform (DXPs). Voyage Journal of Economics & Business Research, 2(2), 93-109.
[5] Dalal, A. (2020). Harnessing the Power of SAP Applications to Optimize Enterprise Resource Planning and Business Analytics. Available at SSRN 5422375.
[6] Hegde, P. (2021). Automated Content Creation in Telecommunications. Jurnal Komputer, Informasi dan Teknologi, 1(2), 20–20.
[7] Dalal, A. (2015). Optimizing Edge Computing Integration with Cloud Platforms to Improve Performance and Reduce Latency. SSRN Electronic Journal. 10.2139/ssrn.5268128.
[8] Bahadur, S., Mondol, K., Mohammad, A., Al-Alam, T., & Bulbul Ahammed, M. (2022). Design and Implementation of Low Cost MPPT Solar Charge Controller.
[9] Dalal, A. (2020). Cyber Threat Intelligence: How to Collect and Analyse Data. International Journal on Recent and Innovation Trends in Computing and Communication.

[10] Mohammad, A., & Mahjabeen, F. (2023). Revolutionizing solar energy: The impact of artificial intelligence on photovoltaic systems. International Journal of Multidisciplinary Sciences and Arts, 2(3), 591856.
[11] Dalal, A. (2023). Data Management Using Cloud Computing. Available at SSRN 5198760.
[12] Dalal, A. (2023). Building Comprehensive Cybersecurity Policies to Protect Sensitive Data in the Digital Era. Available at SSRN 5424094.
[13] Dalal, Aryendra. (2019). Maximizing Business Value through Artificial Intelligence and Machine Learning in SAP Platforms. SSRN Electronic Journal. 10.2139/ssrn.5424315.
[14] Hegde, P. (2019). AI-Powered 5G Networks: Enhancing Speed, Efficiency, and Connectivity. International Journal of Research Science and Management, 6(3), 50-61.
[15] Mohammad, A., Mahjabeen, F., Al-Alam, T., Bahadur, S., & Das, R. (2022). Photovoltaic Power Plants: A Possible Solution for Growing Energy Needs of Remote Bangladesh. Available at SSRN 5185365.
[16] Dalal, A. (2018). Driving Business Transformation through Scalable and Secure Cloud Computing Infrastructure Solutions. Available at SSRN 5424274.
[17] Dalal, A. (2018). Revolutionizing Enterprise Data Management Using SAP HANA for Improved Performance and Scalability. Available at SSRN 5424194.
[18] Dalal, Aryendra. (2022). Addressing Challenges in Cybersecurity Implementation Across Diverse Industrial and Organizational Sectors. SSRN Electronic Journal. 10.2139/ssrn.5422294.
[19] Tiwari, A. (2022). AI-Driven Content Systems: Innovation and Early Adoption. Propel Journal of Academic Research, 2(1), 61–79.
[20] Dalal, A. (2020). Exploring Next-Generation Cybersecurity Tools for Advanced Threat Detection and Incident Response. Available at SSRN 5424096.
[21] Dalal, Aryendra. (2020). Exploring Advanced SAP Modules to Address Industry-Specific Challenges. SSRN Electronic Journal. 10.2139/ssrn.5268100.
[22] Hegde, P., & Varughese, R. J. (2023). Elevating Customer Support Experience in Telecom: AI chatbots, virtual assistants, AR. Propel Journal of Academic Research, 3(2), 193–211.
[23] Tiwari, A. (2023). Generative AI in Digital Content Creation, Curation and Automation. International Journal of Research Science and Management, 10(12), 40–53.
[24] Dalal, A. (2020). Cybersecurity and privacy: Balancing security and individual rights in the digital age. Available at SSRN 5171893.
[25] Dalal, A. (2017). Developing Scalable Applications Through Advanced Serverless Architectures in Cloud Ecosystems. Available at SSRN 5423999.
[26] Maizana, D., Situmorang, C., Satria, H., Yahya, Y. B., Ayyoub, M., Bhalerao, M. V., & Mohammad, A. (2023). The Influence of Hot Point on MTU CB Condition. Journal of Renewable Energy, Electrical, and Computer Engineering, 3(2), 37–43.
[27] Tiwari, A. (2022). Ethical AI Governance in Content Systems. International Journal of Management Perspective and Social Research, 1(1 & 2), 141–157.
[28] Hegde, P., & Varughese, R. J. (2022). Predictive Maintenance in Telecom Using AI. Journal of Mechanical, Civil and Industrial Engineering, 3(3), 102–118.
[29] Dalal, A. (2020). Leveraging Artificial Intelligence to Improve Cybersecurity Defences Against Sophisticated Cyber Threats. Available at SSRN 5422354.
[30] Dalal, Aryendra. (2017). Exploring Emerging Trends in Cloud Computing and Their Impact on Enterprise Innovation. SSRN Electronic Journal. 10.2139/ssrn.5268114.
[31] Dalal, Aryendra. (2018). Leveraging Cloud Computing to Accelerate Digital Transformation Across Diverse Business Ecosystems. SSRN Electronic Journal. 10.2139/ssrn.5268112.
[32] Hegde, P., & Varughese, R. J. (2020). AI-Driven Data Analytics: Insights for Telecom Growth Strategies. International Journal of Research Science and Management, 7(7), 52–68.
[33] Mohammad, A., & Mahjabeen, F. (2023). Promises and challenges of perovskite solar cells: a comprehensive review. BULLET: Jurnal Multidisiplin Ilmu, 2(5), 1147–1157.