Frontiers in Emerging Multidisciplinary Sciences (Open Access)

Volume 03, Issue 8, August 2026,
Publish Date: 13-08-2026
DOI: https://doi.org/10.64917/fems/Volume03Issue08-04 PageNo.36-42

Intelligent Graph Neural Network Framework for Detecting and Classifying Cyber Threats in Cloud Infrastructure

Erion Hoxha
Department of Computer Science, Polytechnic University of Tirana, Albania

Elira Kola
Faculty of Information Technology, University of Tirana, Albania

ABSTRACT

Cloud infrastructure has become increasingly interconnected, dynamic, and heterogeneous, creating security conditions in which cyber threats can propagate across users, virtual machines, containers, services, networks, and data resources. Conventional security mechanisms that primarily analyze isolated events may therefore overlook relational dependencies among entities and activities. This paper proposes an intelligent Graph Neural Network (GNN) framework for detecting and classifying cyber threats in cloud infrastructure by representing cloud security events as a heterogeneous graph and learning threat-related structural and behavioral patterns. The proposed framework integrates graph construction, feature representation, graph-based message propagation, threat classification, and risk-oriented decision generation. Its theoretical foundation is derived from probabilistic networks, dependency modeling, chain-graph representations, causal discovery, and structural learning. The framework is particularly informed by the graph-based deep-learning approach of Marri et al. (2025), which establishes the relevance of graph-based learning for identifying cyber threats in cloud platforms. The methodology conceptualizes cloud entities as nodes and their interactions as edges, allowing the model to capture both local characteristics and higher-order relationships. A multi-stage architecture is developed for preprocessing security telemetry, constructing dynamic graphs, extracting graph representations, identifying anomalous structures, and assigning threat classes. Analytical findings indicate that graph representation can improve contextual threat identification compared with isolated-event analysis because malicious activity is interpreted through relationships among cloud entities. However, challenges involving graph scalability, class imbalance, dynamic topology, weak labels, and explainability remain significant. The proposed framework consequently provides a theoretically grounded architecture for intelligent cloud-threat detection while identifying directions for empirical validation and operational deployment.

KEYWORDS: Graph Neural Networks, Cloud Security, Cyber Threat Detection, Threat Classification, Graph Learning, Cloud Infrastructure, Network Security, Intelligent Detection, Deep Learning.

INTRODUCTION

1.1 Background

Cloud infrastructure operates through tightly coupled computational and communication relationships. Virtual machines exchange traffic, applications access databases, users authenticate with services, containers communicate with other workloads, and security controls continuously generate telemetry. Such interactions form a naturally relational environment in which the meaning of an individual security event can depend strongly on its surrounding entities and activities. Consequently, cyber-threat detection in cloud environments should not be restricted to independent observations when the underlying security problem is inherently relational.

Graph-based modeling provides an appropriate theoretical representation for this environment. A cloud infrastructure can be represented as a graph (G=(V,E)), where (V) represents entities such as users, hosts, workloads, services, IP addresses, and resources, while (E) represents relationships such as authentication, communication, access, process execution, or data transfer. Probabilistic network research demonstrates how dependencies among variables can be represented explicitly rather than treated as unrelated observations (Cowell et al., 2007). Similarly,

multivariate dependency models emphasize that relationships among variables can provide information that is not available from individual variables considered independently (Cox & Wermuth, 2014).

The importance of relational reasoning is also supported by research into chain graphs and structured dependency systems. Andersson et al. (1996) developed an alternative Markov property for chain graphs, demonstrating the value of structured graphical representations for describing conditional relationships. Research on structural learning further establishes that graph structure can encode meaningful dependencies even when some relationships are weak or difficult to observe directly (Eigenmann et al., 2017). These theoretical foundations are relevant to cybersecurity because malicious behavior frequently manifests through sequences or combinations of related activities rather than through a single observable event.

Recent work specifically connecting graph-based deep learning with cloud-security threat identification provides an important technological foundation for the proposed framework. Marri et al. (2025) presented a graph-based deep-learning model for identifying cyber threats in cloud platforms, demonstrating the applicability of graph-oriented learning to cloud-security problems. Building upon this direction, the present study focuses not only on identifying suspicious activity but also on developing an integrated framework for detecting and classifying threats through relational representations.

1.2 Problem Statement

Traditional detection architectures commonly process logs, network flows, authentication records, or system events as individual observations. Such approaches may identify statistical irregularities but can fail to capture relationships among multiple apparently legitimate activities. For example, an isolated login may be normal, an isolated privileged access event may also be normal, and an isolated connection to a service may appear legitimate. However, when these events form a coordinated relationship involving an unusual user, workload, privilege escalation, and subsequent data movement, their collective structure may indicate an attack.

The central research problem is therefore the development of an intelligent model capable of learning security-relevant relationships within cloud infrastructure while simultaneously distinguishing benign and malicious patterns. Marri et al. (2025) provides a direct foundation for graph-based deep-learning detection in cloud platforms, while the broader literature on causal and probabilistic graph structures suggests that relational organization is critical for understanding complex dependencies.

1.3 Research Relevance, Objectives, Scope and Significance

The objective of this research is to formulate an intelligent GNN framework that transforms heterogeneous cloud-security telemetry into graph representations and uses learned relational embeddings for threat detection and classification. The framework aims to capture node-level, edge-level, and neighborhood-level characteristics, enabling security decisions to incorporate contextual relationships.

The scope includes cloud users, virtual machines, containers, applications, services, network endpoints, authentication events, access relationships, and communication patterns. The proposed architecture is applicable conceptually to centralized or distributed cloud environments and is designed to support multiple threat categories rather than binary anomaly detection alone.

The significance of the research lies in combining graph-theoretic dependency modeling with deep representation learning. Rather than treating cloud security as an exclusively feature-vector classification problem, the framework positions security analysis as a structured relational-learning problem.

2. LITERATURE REVIEW

Graphical models provide an important theoretical basis for representing dependencies among interconnected variables. Cowell et al. (2007) established probabilistic networks as formal mechanisms for representing and computationally analyzing dependencies, while Cox and Wermuth (2014) emphasized multivariate dependency structures and their interpretation. For cloud cybersecurity, these principles imply that security states can be modeled through relationships among infrastructure entities rather than only through independent feature values.

Andersson et al. (1996) investigated Markov properties for chain graphs, illustrating how graphical structures can represent conditional relationships under more general configurations than conventional directed or undirected graphs. This perspective is useful for cloud environments because interactions may include directional communication, shared dependencies, and partially symmetric relationships.

Research on causal discovery strengthens the theoretical argument for structured reasoning. Andrews et al. (2020) examined causal discovery under latent confounding and tiered background knowledge, demonstrating that

unobserved variables and prior structural knowledge influence the interpretation of graph relationships. In cybersecurity, hidden or unrecorded factors can similarly affect observed security events. Therefore, a threat-detection graph should not automatically interpret every observed association as a direct causal relationship.

Bhattacharya et al. (2020) considered causal inference under interference and network uncertainty. This is particularly relevant to cloud infrastructure because the behavior of one entity can affect or alter the security state of connected entities. A compromised workload, for example, can influence neighboring services through authentication, communication, or shared-resource relationships.

Eigenmann et al. (2017) addressed structure learning for linear Gaussian structural equation models with weak edges. Weak relationships are significant in security analytics because low-frequency interactions may initially appear insignificant but can become important when combined with stronger relationships. A graph-based threat detector should therefore avoid eliminating weak connections solely because they have low individual statistical strength.

The biological network studies of Balaji et al. (2006) and Gama-Castro et al. (2008) provide additional evidence for the analytical usefulness of network representations. Although these studies concern regulatory systems rather than cybersecurity, they demonstrate how complex systems can be represented through interacting entities and regulatory relationships. Their methodological implication for cloud security is that network-level structure can reveal system behavior that cannot be adequately represented by isolated variables.

Chen et al. (2018) demonstrated the integration of semi-supervised learning and graph-based factorization for recommendation. This provides an important conceptual insight for cybersecurity, where obtaining high-quality labels for every security event is difficult. Semi-supervised graph learning can potentially exploit both labeled attack examples and unlabeled operational activity.

Collectively, the literature establishes three major foundations: graphical representations capture dependencies; structural and causal learning address complex relationships and uncertainty; and machine-learning approaches can exploit graph structures for predictive tasks. The direct cloud-security contribution of Marri et al. (2025) further supports the feasibility of graph-based deep learning for identifying threats in cloud platforms. However, an important research gap remains in integrating heterogeneous cloud entities, dynamic relationships, threat detection, and multi-class classification within a unified intelligent architecture. The proposed framework addresses this gap conceptually.

3. METHODOLOGY

3.1 Proposed Framework Architecture

The proposed framework consists of six interconnected stages: cloud telemetry acquisition, preprocessing and normalization, dynamic graph construction, graph representation learning, threat detection and classification, and risk-oriented decision generation.

At the first stage, security telemetry is collected from authentication systems, network flows, workload activity, access-control mechanisms, application logs, and infrastructure monitoring systems. Each observation is converted into a structured security event containing attributes such as timestamp, source entity, destination entity, action type, privilege level, protocol, resource identifier, and behavioral indicators.

The second stage performs preprocessing. Missing values are handled through appropriate imputation or event exclusion, categorical attributes are encoded, numerical variables are normalized, and timestamps are transformed into temporal features. Events with inconsistent identifiers are resolved so that the same cloud entity is represented consistently across the graph.

3.2 Dynamic Cloud Security Graph

The preprocessed telemetry is converted into a heterogeneous graph:

[

G_t=(V_t,E_t,X_t)

]

where (V_t) represents cloud entities at time (t), (E_t) represents interactions, and (X_t) contains node and edge attributes.

Nodes may represent users, virtual machines, containers, applications, services, databases, storage resources, and network endpoints. Edges represent relationships including authentication, network communication, process interaction, privilege assignment, API invocation, and resource access.

A dynamic representation is preferred because cloud infrastructure changes continuously. Nodes can appear or disappear as workloads are created or terminated, while edges change according to operational behavior. This design

also reflects the network-uncertainty considerations discussed by Bhattacharya et al. (2020).

3.3 Graph Feature Representation

Node features represent intrinsic characteristics, such as authentication frequency, average traffic volume, privilege level, failed-login ratio, resource-access diversity, and workload behavior. Edge features describe relationship characteristics, including communication frequency, protocol, duration, data volume, and access type.

The framework distinguishes three levels of information. Node-level features describe individual entities; edge-level features describe interactions; and neighborhood-level features describe the surrounding behavioral context. This distinction is critical because a suspicious node may not be identifiable from its intrinsic attributes alone.

3.4 Graph Neural Network Learning

The GNN performs iterative neighborhood aggregation. A generic layer can be expressed as:

[

h_v^{(l+1)}=\sigma \left(W_1h_v^{(l)}+

\sum_{u\in N(v)}\alpha_{vu}W_2h_u^{(l)}\right)

]

where (h_v^{(l)}) represents the embedding of node (v) at layer (l), (N(v)) is its neighborhood, (W_1) and (W_2) are learnable parameters, (\alpha_{vu}) represents the contribution of neighboring node (u), and (\sigma) is an activation function.

Through repeated aggregation, the representation of an entity incorporates increasingly broader contextual information. This mechanism is particularly appropriate for cyberattack detection because attack behavior may propagate through several connected entities. The graph-based deep-learning direction demonstrated by Marri et al. (2025) provides the principal methodological motivation for applying such representations to cloud threat identification.

3.5 Detection and Classification Layer

The learned graph embeddings are supplied to two complementary decision functions. The first estimates whether a graph entity or subgraph exhibits suspicious behavior. The second assigns the detected behavior to a threat category.

A softmax classification function can be expressed as:

[

P(y=k|h_v)=

\frac{\exp(z_k)}

{\sum_{j=1}^{K}\exp(z_j)}

]

where (K) is the number of threat classes and (z_k) is the classification score for class (k).

Potential classes include benign activity and multiple malicious categories such as credential misuse, anomalous access, lateral movement, denial-oriented activity, or suspicious data transfer. The exact class taxonomy should be defined according to the operational dataset used for empirical validation.

3.6 Semi-Supervised Learning Strategy

Because cybersecurity datasets commonly contain far more benign events than confirmed attack events, the framework supports semi-supervised learning. A limited set of labeled nodes or subgraphs can provide attack supervision while unlabeled graph structures contribute contextual information.

This design is theoretically compatible with the graph-based semi-supervised perspective presented by Chen et al. (2018). It also addresses the practical difficulty of obtaining comprehensive labels for every cloud-security event.

3.7 Causal and Structural Considerations

The framework does not interpret graph connectivity as automatic causality. Andrews et al. (2020) demonstrates that latent confounding and background knowledge can affect causal interpretation, while Eigenmann et al. (2017) shows the relevance of weak relationships in structural learning. Therefore, graph edges should be treated primarily as observed or engineered relationships unless causal evidence is separately established.

A further consideration is conditional dependency. The graphical modeling principles discussed by Andersson et al. (1996) and Cowell et al. (2007) suggest that security relationships can be interpreted through conditional structures rather than through simple pairwise associations.

3.8 Threat Decision and Alert Generation

The final stage converts classification probabilities and graph-context information into a security risk score. A conceptual risk function is:

[

R_v=\lambda_1P_a+\lambda_2S_v+\lambda_3C_v

]

where (P_a) denotes predicted attack probability, (S_v) represents the severity associated with the predicted threat, (C_v) represents contextual criticality, and (\lambda_1,\lambda_2,\lambda_3) are weighting parameters.

This enables security systems to prioritize high-confidence threats affecting critical infrastructure rather than generating identical alerts for all anomalous events.

4. RESULTS

The analytical findings of the proposed framework indicate that graph-based representation provides a stronger conceptual basis for cloud-threat analysis than treating security events as independent records. The primary finding is that threat evidence can be distributed across multiple related entities. A suspicious authentication event, for example, may have limited significance in isolation but become substantially more informative when connected to an unusual workload, privilege change, and subsequent communication with a sensitive service. The GNN architecture captures this contextual information through neighborhood aggregation.

A second finding concerns threat classification. Conventional binary anomaly detection can indicate that an observation is unusual without explaining the likely threat category. The proposed classification layer transforms learned graph representations into multi-class threat probabilities, thereby supporting differentiated security responses. This extends the graph-based threat-identification direction reported by Marri et al. (2025) toward a unified detection-and-classification architecture.

The third finding is the importance of heterogeneous relationships. Cloud infrastructure does not contain only network connections; it also contains identity, access-control, process, application, and resource relationships. Representing these relationships within a heterogeneous graph allows the model to associate behavioral changes across infrastructure layers. Such relational reasoning is consistent with the broader dependency-modeling principles discussed by Cowell et al. (2007) and Cox and Wermuth (2014).

A fourth finding is that weak and indirect relationships should not automatically be discarded. Eigenmann et al. (2017) demonstrates the relevance of weak edges in structural learning, while Bhattacharya et al. (2020) highlights network uncertainty and interference. In cloud security, an individually weak interaction may become significant when it participates in a coordinated attack path. Therefore, edge weighting and neighborhood aggregation should preserve potentially informative low-strength relationships while controlling noise.

The framework also indicates the value of semi-supervised learning. Complete labeling of cloud telemetry is operationally expensive, whereas unlabeled events are abundant. The semi-supervised perspective of Chen et al. (2018) therefore provides a useful methodological foundation for learning from limited attack labels combined with large volumes of unlabeled infrastructure activity.

However, the findings also identify several limitations. Graph construction itself can become computationally expensive in large cloud environments. Dynamic topology can cause frequent changes in node and edge representations, while class imbalance may cause the model to favor benign activity. Furthermore, graph connectivity does not necessarily imply causation. The causal-discovery literature, particularly Andrews et al. (2020), indicates that latent variables and background assumptions can alter causal interpretations. Consequently, the proposed architecture should be understood as a predictive detection framework rather than an automatic causal-analysis system.

Overall, the findings suggest that graph-based deep learning is particularly valuable when security evidence is distributed across interconnected cloud entities. The framework provides a mechanism for integrating local behavior, relational context, and classification probabilities into a unified security decision process.

5. DISCUSSION

The proposed framework contributes theoretically by positioning cloud-threat detection as a structured relational-learning problem. Traditional feature-based classification assumes that the principal information needed for prediction is contained within each individual observation. In contrast, the proposed GNN assumes that security meaning is partly encoded in relationships among observations. This distinction is important because cloud attacks frequently involve coordinated activity across identities, workloads, services, and network resources. The graph-based approach presented by Marri et al. (2025) provides direct support for this conceptual shift.

The framework also extends the implications of probabilistic and dependency-based network research.

Cowell et al. (2007) and Cox and Wermuth (2014) demonstrate that structured dependencies can provide a richer representation of complex systems than isolated variables. Applying this principle to cloud infrastructure allows the proposed model to integrate multiple forms of security evidence. The architecture consequently has potential to identify attack patterns whose distinguishing characteristic is not an extreme value in one feature but an unusual configuration of several connected events.

An important theoretical trade-off concerns predictive learning versus causal interpretation. Andrews et al. (2020) shows that latent confounding can complicate causal discovery, while Bhattacharya et al. (2020) demonstrates that network uncertainty and interference create additional challenges. Accordingly, a high-probability GNN prediction should not automatically be interpreted as evidence that one cloud event caused another. The framework should instead provide predictive evidence and, where required, support subsequent investigation by security analysts.

The literature on chain graphs and structural models also provides a useful interpretive foundation. Andersson et al. (1996) demonstrates how conditional relationships can be represented through graphical structures, while Eigenmann et al. (2017) highlights the importance of weak edges. These insights suggest that cloud-security graphs should preserve structural information rather than aggressively simplify the network. Nevertheless, retaining every relationship can increase computational complexity and introduce noise. Practical implementations therefore require graph-pruning or edge-weighting strategies that balance information retention against scalability.

The framework's semi-supervised capability represents another important practical implication. Chen et al. (2018) illustrates how graph-based learning can exploit partially labeled data. For cloud-security operations, this is significant because confirmed attack labels are generally more difficult to obtain than ordinary operational records. A model capable of learning from both labeled and unlabeled structures could reduce dependence on exhaustive manual annotation.

The network-oriented studies of Balaji et al. (2006) and Gama-Castro et al. (2008) further support the broader proposition that system behavior can be understood through interaction networks. Although their application domain differs substantially from cybersecurity, their network-centric analytical logic is transferable at the conceptual level. The implication is not that biological regulatory networks and cloud infrastructure are equivalent, but that relational structure can serve as a meaningful analytical abstraction across complex systems.

Despite these advantages, several limitations must be addressed before operational deployment. First, empirical validation on representative cloud datasets is required to determine detection accuracy, precision, recall, F1-score, false-positive rate, computational latency, and scalability. Second, dynamic graph updates must be engineered carefully to prevent outdated relationships from distorting current security decisions. Third, explainability remains important because security analysts need understandable evidence for why a graph region was classified as malicious. Finally, adversarial manipulation of graph structure itself represents a potential weakness that future research should investigate.

Thus, the framework should be regarded as an integrated research architecture rather than a claim of universal detection superiority. Its primary contribution is the systematic combination of graph representation, deep relational learning, semi-supervised modeling, multi-class threat classification, and contextual risk prioritization for cloud infrastructure.

6. CONCLUSION

This paper proposed an intelligent Graph Neural Network framework for detecting and classifying cyber threats in cloud infrastructure. The framework represents cloud users, workloads, services, resources, and security events as interconnected graph entities and applies relational representation learning to identify suspicious patterns. Its theoretical foundation integrates probabilistic networks, dependency modeling, structural learning, causal reasoning, and semi-supervised graph learning.

The principal contribution is the integration of detection and multi-class classification within a dynamic cloud-security graph. Rather than evaluating security events independently, the framework incorporates neighborhood and interaction context, enabling coordinated behavioral patterns to influence threat decisions. The approach is theoretically supported by graph-based cloud-threat research reported by Marri et al. (2025) and by broader studies of structured networks and dependencies.

Future research should empirically validate the architecture using real-world cloud-security datasets, investigate temporal and heterogeneous GNN variants, address class imbalance, develop explainable threat paths, and evaluate resilience against adversarial graph manipulation. Such extensions could establish whether graph-based deep learning can provide scalable, accurate, and operationally

interpretable threat intelligence for increasingly complex cloud infrastructures.

REFERENCES

1. Andersson, S. A., Madigan, D., & Perlman, M. D. (1996). An alternative markov property for chain graphs. In Proceedings of the 12th Conference on Uncertainty in Artificial Intelligence, pp. 40–48. AI.
2. Andrews, B., Spirtes, P., & Cooper, G. F. (2020). On the completeness of causal discovery in the presence of latent confounding with tiered background knowledge. In Proceedings of the 23th International Conference on Artificial Intelligence and Statistics, pp. 4002–4011. PMLR.
3. Balaji, S., Babu, M. M., Iyer, L. M., Luscombe, N. M., & Aravind, L. (2006). Comprehensive analysis of combinatorial regulation using the transcriptional regulatory network of yeast. Journal of molecular biology, 360(1), 213–227.
4. Bhattacharya, R., Malinsky, D., & Shpitser, I. (2020). Causal inference under interference and network uncertainty. In Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence, pp. 1028–1038. PMLR.
5. Chen, C., Chang, K. C.-C., Li, Q., & Zheng, X. (2018). Semi-supervised learning meets factorization: Learning to recommend with chain graph model. ACM Transactions on Knowledge Discovery from Data, 12(6), 1–24.
6. Cowell, R. G., Dawid, P., Lauritzen, S. L., & Spiegelhalter, D. J. (2007). Probabilistic networks and expert systems: Exact computational methods for Bayesian networks. Springer Science & Business Media.
7. Cox, D. R., & Wermuth, N. (2014). Multivariate dependencies: Models, analysis and interpretation. Chapman and Hall/CRC.
8. Eigenmann, M. F., Nandy, P., & Maathuis, M. H. (2017). Structure learning of linear gaussian structural equation models with weak edges. In Proceedings of the 33th Conference on Uncertainty in Artificial Intelligence.
9. Gama-Castro, S., Jiḿenez-Jacinto, V., Peralta-Gil, M., Santos-Zavaleta, A., Pẽnaloza-Spinola, M. I., Contreras-Moreira, B., Segura-Salazar, J., Muniz-Rascado, L., Martinez-Flores, I., Salgado, H., et al. (2008). Regulondb (version 6.0): gene regulation model of escherichia coli k-12 beyond transcription, active (experimental) annotated promoters and textpresso navigation. Nucleic acids research, 36(suppl1), D120–D124.
10. M. R. Marri, S. B. Kurada, S. Gupta, S. Dey, S. Kumar and R. Chauhan, "Graph-Based Deep Learning Model for Identifying Cyber Threats in Cloud Platforms," 2025 International Conference on Computational Intelligence, Security, and Artificial Intelligence (IntelliSecAI), Al-Khobar, Saudi Arabia, 2025, pp. 1-7, doi: 10.1109/IntelliSecAI66368.2025.11472831.