



Journal Website:
<http://sciencebring.com/index.php/ijasr>

Copyright: Original content from this work may be used under the terms of the creative commons attributes 4.0 licence.

Research Article

Graph Deep Learning-Driven Cyber Threat Detection and Risk Analysis Framework for Secure

Submission Date: June 29, 2026, **Accepted Date:** July 18, 2026,

Published Date: August 13, 2026

Crossref doi: <https://doi.org/10.37547/ijasr-06-08-07>

Abdullah Al-Harbi

College of Computer and Information Sciences, King Saud University, Saudi Arabia

Reem Al-Qahtani

Department of Artificial Intelligence, Princess Nourah bint Abdulrahman University, Saudi Arabia

ABSTRACT

The increasing complexity of cloud infrastructures has created security environments in which cyber threats emerge from interconnected relationships among users, applications, virtual machines, services, network flows, and access privileges. Conventional security analytics frequently represent these entities as independent observations, limiting their ability to identify relational attack patterns and propagate contextual risk. This paper proposes a Graph Deep Learning-Driven Cyber Threat Detection and Risk Analysis Framework for Secure Cloud Environments that models cloud infrastructures as dynamic heterogeneous graphs and applies graph-based representation learning to identify anomalous entities, suspicious relationships, and coordinated attack behaviors. The framework integrates graph construction, feature representation, graph deep learning, threat classification, temporal reasoning, and risk prioritization into a unified security pipeline. Its theoretical foundation is strengthened by research on symbolic representations, state-space abstraction, generalized planning, temporal logic, and structured reasoning. In particular, concepts of relational representation and structural abstraction provide a basis for transforming complex cloud security states into machine-processable graph structures. The proposed framework extends graph-based cyber threat identification toward risk-aware decision support by associating detected threats with affected assets, attack relationships, and potential propagation paths. The analysis indicates that graph deep learning can provide greater contextual awareness than isolated event classification, while temporal and symbolic reasoning can improve interpretability and decision

consistency. The framework is particularly relevant to cloud environments characterized by dynamic workloads, distributed services, privilege dependencies, and rapidly changing attack surfaces.

KEYWORDS

Graph Deep Learning, Cyber Threat Detection, Cloud Security, Graph Neural Networks, Risk Analysis, Threat Classification, Temporal Reasoning, Security Analytics, Cloud Infrastructure

INTRODUCTION

1.1 Background

Cloud computing has transformed infrastructure security from a relatively static perimeter-oriented problem into a highly interconnected and continuously changing analytical problem. A cloud environment may simultaneously contain users, identities, containers, virtual machines, storage resources, applications, APIs, network connections, authentication events, and privilege relationships. Consequently, an observed security event cannot always be interpreted correctly without considering its relationships with other entities.

Graph-based learning provides a natural representation for such environments because nodes can represent cloud entities while edges can describe communication, authentication, ownership, dependency, or access relationships. The central premise of the proposed approach is that cyber threats should therefore be analyzed not only as anomalous individual events but also as abnormal configurations or evolving patterns within an infrastructure graph. Marri et al. (2025) demonstrate the relevance of graph-based deep learning for identifying cyber threats in cloud

platforms, providing an important foundation for extending graph representation toward integrated threat detection and risk analysis.

The theoretical basis for structured representation can also be related to work in artificial intelligence concerning state-space representation, abstraction, relational features, and generalized planning. Bonet and Geffner (2018) emphasize the importance of features and projections for changing representations while retaining information relevant to planning. Similarly, Bonet and Geffner (2020a) investigate learning first-order symbolic representations from the structure of state spaces. These perspectives are valuable for cybersecurity because cloud infrastructures can be interpreted as structured states whose entities and relationships evolve over time.

1.2 Problem Statement

Many security analytics approaches treat logs, flows, authentication events, and resource activities as independent records. Such representations may identify unusual behavior but can struggle to explain whether several individually benign activities constitute a

coordinated threat. For example, an unusual login, followed by privilege escalation, access to a sensitive storage resource, and communication with another cloud service may collectively indicate an attack even when none of the events is sufficiently anomalous in isolation.

The problem is therefore not limited to classification accuracy. A practical cloud security system must identify suspicious graph structures, assess the severity of detected threats, determine affected assets, and support prioritization. Furthermore, cloud security states are temporal: the order and duration of activities may determine whether a relationship represents legitimate behavior or an attack sequence. Temporal logic research demonstrates how temporal constraints can be used to represent search-control knowledge and action sequences (Bacchus & Kabanza, 2000). This principle can inform temporal threat analysis in cloud graphs.

1.3 Research Objectives

This research develops a conceptual graph deep learning framework with five objectives: to represent cloud infrastructures as dynamic heterogeneous graphs; to learn relational representations for threat detection; to classify suspicious nodes, edges, and subgraphs; to incorporate temporal and structural reasoning into risk analysis; and to generate risk-prioritized security outputs.

1.4 Scope and Significance

The proposed framework focuses on cloud infrastructure security rather than a specific

cloud provider. It addresses identity activities, resource interactions, network communication, access relationships, and security events. Its significance lies in integrating detection with contextual risk analysis. Rather than merely reporting an anomaly, the framework seeks to determine its position within the broader infrastructure and estimate its potential impact.

LITERATURE REVIEW

Existing literature supplied for this study spans symbolic reasoning, planning, representation learning, temporal logic, and graph-based cyber threat detection. Although most of the earlier works are not specifically designed for cybersecurity, their theoretical contributions are directly relevant to constructing structured security representations.

Baader et al. (2003) provide a formal foundation for description logics, demonstrating how structured entities and relationships can be represented using formal knowledge structures. For cloud security, this perspective supports modeling users, resources, permissions, and dependencies as semantically related objects. Bacchus (2001) illustrates the importance of systematic evaluation and representation in planning systems, while Bacchus and Kabanza (2000) demonstrate how temporal logic can encode constraints over action sequences. The latter is particularly relevant when security events must be evaluated according to their temporal ordering.

Baier et al. (2008) move beyond classical planning by emphasizing procedural control knowledge and preferences. Their work suggests that decision systems can benefit from mechanisms that distinguish not only what actions are possible but also which actions are preferable under particular conditions. In cybersecurity, this idea can be transferred to threat prioritization, where multiple detected anomalies may require different levels of response.

Asai (2019) investigates unsupervised grounding of plannable first-order representations from images. Although the application domain differs substantially, the underlying concept of learning structured representations from observations provides useful theoretical motivation for automatically converting raw cloud telemetry into higher-level graph structures.

Bonet and Geffner (2018) examine features, projections, and representation changes for generalized planning. Their emphasis on abstraction is important because cloud environments can produce extremely large security graphs. A practical detection system must retain security-relevant structure while reducing irrelevant complexity. Bonet et al. (2019) further explore learning features and abstract actions for generalized plans, suggesting that learned abstractions can support reasoning across different states.

Bonet and Geffner (2020a) focus on first-order symbolic representations learned from state-space structure, while Bonet and Geffner (2020b)

investigate qualitative numeric planning and its reductions. Together, these studies reinforce the idea that structured representations can improve reasoning over complex environments. Bonet and Geffner (2021) further examine general policies, representations, and planning width, providing a conceptual basis for designing scalable representations for large state spaces.

Daggelinckx (2023) addresses planning sketches and verification through temporal logic, reinforcing the relevance of combining abstraction with temporal validation. Dalmau-Moreno et al. (2023) investigate combined task and motion planning through sketch decompositions, illustrating how complex problems can be decomposed into structured subproblems.

Within the specific cybersecurity domain, Marri et al. (2025) provide the most direct foundation for the present study by applying graph-based deep learning to cyber threat identification in cloud platforms. Building on this contribution, the present framework extends the emphasis from threat identification toward integrated detection, contextual reasoning, and risk analysis.

The principal research gap is therefore the limited integration of graph-based threat detection with structured representation, temporal reasoning, and explicit risk prioritization. Existing planning and representation studies offer theoretical mechanisms for abstraction and temporal reasoning, while the cloud cybersecurity contribution establishes the relevance of graph

deep learning. The proposed framework connects these perspectives into a unified security architecture.

METHODOLOGY

3.1 Research Design

This study adopts a framework-development methodology. Rather than claiming results from an unperformed experimental deployment, it defines an analytically grounded architecture that can subsequently be implemented and empirically evaluated. The methodology consists of six stages: cloud telemetry acquisition, heterogeneous graph construction, graph representation learning, threat detection and classification, temporal-risk analysis, and security decision generation.

3.2 Dynamic Cloud Security Graph

The cloud environment is represented as a time-dependent graph:

$$G_t = (V_t, E_t, X_t)$$

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

Nodes may include users, identities, virtual machines, containers, databases, applications, APIs, storage objects, and network endpoints. Edges represent interactions such as login, access,

communication, invocation, privilege assignment, or resource dependency.

A heterogeneous graph is preferable because different node and relationship types carry different security meanings. For example, a user-to-resource access edge differs semantically from a virtual-machine-to-network communication edge. The framework therefore assigns relation-specific representations rather than treating all connections identically.

3.3 Feature Representation

Each graph entity is transformed into a feature vector containing behavioral, structural, and contextual information. Potential features include authentication frequency, access frequency, connection patterns, privilege level, resource sensitivity, historical activity, and neighborhood characteristics.

The representation process follows the principle that useful abstractions should preserve information relevant to downstream reasoning. This is consistent with the representation-change perspective discussed by Bonet and Geffner (2018). For large cloud graphs, graph sampling and hierarchical abstraction can reduce computational complexity while retaining suspicious relationships.

3.4 Graph Deep Learning Layer

The graph learning component generates contextual embeddings by aggregating information from neighboring nodes and edges.

Conceptually, the representation of node (v) at layer ($l+1$) can be expressed as:

$$[h_v^{(l+1)} = \sigma(W_{0h_v^{(l)}} + \sum_{u \in N(v)} \alpha_{uv} W_{rh_u^{(l)}})]$$

where ($N(v)$) denotes neighboring entities, (W_r) represents relation-specific transformations, (α_{uv}) represents learned relational importance, and (σ) denotes a nonlinear transformation.

This mechanism allows the model to distinguish isolated anomalies from anomalies embedded in suspicious neighborhoods. For example, a normally functioning user may receive a high-risk representation if its activity becomes connected to multiple compromised resources.

The approach extends the graph-learning direction established by Marri et al. (2025) by incorporating graph embeddings into a broader risk-analysis pipeline rather than limiting the output to threat identification.

3.5 Threat Detection and Classification

The framework produces three complementary outputs: anomaly score, threat category, and graph-level threat confidence. The anomaly score captures deviation from learned behavioral patterns, while classification determines the likely threat category. Graph-level confidence evaluates whether multiple suspicious activities form a coherent attack structure.

A hypothetical example is an identity that normally accesses one application but suddenly authenticates from an unusual context, accesses a privileged service, communicates with a sensitive database, and subsequently initiates multiple resource operations. The framework evaluates these events jointly rather than treating them as unrelated records.

3.6 Temporal and Structural Risk Analysis

Temporal reasoning is incorporated to preserve event order. A sequence such as authentication → privilege change → sensitive-resource access → abnormal communication is potentially more significant than the same events occurring independently. Temporal logic provides an established theoretical foundation for expressing constraints over ordered activities (Bacchus & Kabanza, 2000).

The risk score can be represented conceptually as:

$$[R = \lambda_1 T + \lambda_2 A + \lambda_3 P + \lambda_4 S + \lambda_5 C]$$

where (T) represents threat confidence, (A) affected-asset criticality, (P) privilege exposure, (S) structural propagation potential, and (C) temporal coherence. The coefficients (λ_i) can be calibrated during implementation.

This formulation allows the same detected threat to receive different priorities depending on its

context. An anomaly involving a low-value test resource may receive a lower risk score than a similar anomaly involving a highly privileged identity and sensitive database.

3.7 Risk Prioritization and Decision Layer

The final layer converts analytical outputs into ranked security alerts. This reduces the limitation of conventional anomaly detection systems that can produce large numbers of alerts without contextual prioritization. The decision layer can categorize threats into low, moderate, high, and critical levels.

The concept of preferences in planning research (Baier et al., 2008) provides theoretical support for this prioritization principle. Instead of considering all security responses equivalent, the framework ranks alerts according to their expected impact and structural significance.

3.8 Verification and Explainability

The framework incorporates structural explanations by identifying influential nodes, relationships, and temporal sequences contributing to a risk score. This is important because graph models may otherwise provide predictions without sufficient operational justification.

The verification concept is consistent with Daggelinckx (2023), where temporal logic is used for verifying planning sketches. In the proposed security setting, verification can determine whether a detected graph pattern satisfies

predefined temporal and structural conditions before escalating it.

RESULTS

The proposed framework produces several analytically significant findings. First, graph representation provides a stronger contextual basis for cloud threat detection than isolated-event representations. Cloud activities are inherently relational: authentication connects identities with resources, network flows connect workloads, and privilege relationships connect users with administrative capabilities. Modeling these dependencies allows the detection system to identify suspicious combinations rather than relying exclusively on individual-event deviations. This direction is consistent with the graph-based cloud threat identification approach presented by Marri et al. (2025).

Second, the framework indicates that graph representation should be hierarchical. A complete cloud graph may contain millions of entities and relationships, making unrestricted processing inefficient. Representation abstraction, motivated by Bonet and Geffner (2018, 2020a), provides a mechanism for preserving security-relevant structural information while reducing unnecessary computational complexity. The implication is that scalable threat detection requires both expressive graph learning and selective structural compression.

Third, temporal reasoning materially changes risk interpretation. Two graph structures with similar entities and relationships may represent

different levels of danger when their event sequences differ. The integration of temporal constraints therefore provides a mechanism for distinguishing legitimate complex workflows from coordinated attacks. This finding is theoretically compatible with temporal planning research by Bacchus and Kabanza (2000) and temporal verification concepts discussed by Daggelinckx (2023).

Fourth, risk analysis benefits from combining detection confidence with asset criticality and structural propagation. A highly confident anomaly does not necessarily constitute the highest operational priority if it affects an isolated, low-value resource. Conversely, a moderate anomaly connected to privileged identities and sensitive resources may warrant immediate investigation. The proposed risk function therefore shifts security analytics from binary detection toward contextual prioritization.

Fifth, structured representation contributes to interpretability. By tracing the subgraph responsible for an alert, analysts can examine which identity, resource, relationship, or sequence contributed to the decision. This is particularly important in cloud environments where opaque classifications may be difficult to operationalize. The theoretical emphasis on structured representations and abstractions across the planning literature supports this design principle (Bonet & Geffner, 2021).

Finally, the framework reveals an important trade-off between model complexity and operational scalability. More expressive

heterogeneous and temporal graphs can capture richer threat relationships but increase computational and implementation requirements. Consequently, the proposed architecture is best interpreted as a modular framework in which graph resolution, temporal depth, and feature complexity can be adjusted according to infrastructure size and security requirements.

DISCUSSION

The principal contribution of the proposed framework is its integration of graph deep learning with structured and temporal risk analysis. Existing planning literature demonstrates that representation determines the quality of reasoning possible over complex states. Bonet and Geffner (2018) and Bonet and Geffner (2020a) particularly emphasize representation transformation and learned symbolic structures, concepts that translate naturally to cloud security graphs. A cloud infrastructure can be interpreted as a dynamic state in which entities and relationships change as users authenticate, workloads communicate, permissions evolve, and resources are created or removed.

The framework also extends the practical direction of Marri et al. (2025). Their graph-based deep learning model establishes a direct connection between graph learning and cyber threat identification in cloud platforms. The present architecture broadens this concept by introducing an explicit risk-analysis layer. This distinction is important because detection alone

does not resolve the operational problem of alert prioritization. A security system must determine not only whether behavior is suspicious but also why it matters and what infrastructure context increases its potential impact.

Temporal reasoning represents another significant theoretical extension. Bacchus and Kabanza (2000) show that temporal logic can express constraints over actions and their ordering. In cloud security, attack progression similarly depends on temporal relationships. A sequence of low-severity events may become highly significant when the events form a recognizable progression toward privilege escalation or sensitive-resource access. Therefore, temporal constraints can reduce false positives generated by evaluating each event independently.

Nevertheless, several limitations remain. First, graph construction quality directly affects model quality. Incorrect or incomplete relationships can produce misleading embeddings. Second, dynamic cloud environments require continuous graph updates, creating computational and data-management challenges. Third, deep graph models may remain difficult to interpret despite structural explanation mechanisms. Fourth, risk coefficients require careful calibration because arbitrary weighting can distort prioritization.

The theoretical literature also suggests that abstraction introduces a trade-off. Bonet et al. (2019) and Bonet and Geffner (2021) demonstrate the value of abstract representations for generalized reasoning, but

excessive abstraction can eliminate distinctions necessary for accurate threat identification. Security graph abstraction must therefore preserve critical attributes such as privilege, resource sensitivity, temporal order, and relationship type.

The framework consequently should not be viewed as a replacement for established security controls. Instead, it provides an intelligent analytical layer that can correlate heterogeneous evidence and prioritize investigation. Future empirical implementation should compare the architecture against conventional machine-learning classifiers and graph-learning baselines using precision, recall, F1-score, false-positive rate, detection latency, and risk-ranking effectiveness.

CONCLUSION

This paper proposed a Graph Deep Learning-Driven Cyber Threat Detection and Risk Analysis Framework for Secure Cloud Environments. The framework models cloud infrastructure as a dynamic heterogeneous graph and combines graph representation learning, threat classification, temporal reasoning, structural analysis, and risk prioritization.

The analysis demonstrates that graph deep learning can provide contextual threat representations by learning relationships among users, workloads, resources, and security events. The theoretical contributions of structured representation, abstraction, temporal logic, and generalized reasoning provide a foundation for

extending graph-based cybersecurity analysis beyond isolated anomaly detection. The framework also builds upon the cloud-focused graph threat detection direction of Marri et al. (2025) by incorporating explicit risk assessment and alert prioritization.

Its principal contribution is therefore conceptual integration: detection identifies suspicious behavior, graph structure explains its context, temporal reasoning evaluates its progression, and risk analysis determines its operational significance. Future research should implement the framework on real cloud telemetry, investigate heterogeneous graph architectures, evaluate temporal graph learning, calibrate risk functions, and measure performance under evolving attack scenarios. Such empirical validation would establish whether the proposed integration can provide measurable improvements in detection accuracy, scalability, explainability, and security-response prioritization.

REFERENCES

1. Asai, M. (2019). Unsupervised grounding of plannable first-order logic representation from images. In Lipovetzky, N., Onaindia, E., & Smith, D. E. (Eds.), *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling (ICAPS 2019)*, pp. 583–591. AAAI Press.
2. Baader, F., Calvanese, D., McGuinness, D. L., Nardi, D., & Patel-Schneider, P. F. (Eds.). (2003). *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press.
3. Bacchus, F. (2001). The AIPS'00 planning competition. *AI Magazine*, 22(3), 47–56.
4. Bacchus, F., & Kabanza, F. (2000). Using temporal logics to express search control knowledge for planning. *Artificial Intelligence*, 116(1–2), 123–191.
5. Baier, J. A., Fritz, C., Bienvenu, M., & McIlraith, S. A. (2008). Beyond classical planning: Procedural control knowledge and preferences in state-of-the-art planners. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI 2008)*, pp. 1509–1512. AAAI Press.
6. Bonet, B., Francès, G., & Geffner, H. (2019). Learning features and abstract actions for computing generalized plans. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI 2019)*, pp. 2703–2710. AAAI Press.
7. Bonet, B., & Geffner, H. (2018). Features, projections, and representation change for generalized planning. In Lang, J. (Ed.), *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI 2018)*, pp. 4667–4673. IJCAI.
8. Bonet, B., & Geffner, H. (2020a). Learning first-order symbolic representations for planning from the structure of the state space. In De Giacomo, G. (Ed.), *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI 2020)*, pp. 2322–2329. IOS Press.
9. Bonet, B., & Geffner, H. (2020b). Qualitative numeric planning: Reductions and

complexity. Journal of Artificial Intelligence Research, 69, 923–961.

10. Bonet, B., & Geffner, H. (2021). General policies, representations, and planning width. In Leyton-Brown, K., & Mausam (Eds.), Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2021), pp. 11764–11773. AAAI Press.
11. Daggelinckx, A. (2023). Generating and verifying planning sketches using temporal logic. Master's thesis, Utrecht University.
12. Dalmau-Moreno, M., García, N., Gómez, V., & Geffner, H. (2023). Combined task and motion planning via sketch decompositions. In ICAPS 2023 Planning and Robotics Workshop (PlanRob).
13. 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.