COLOMBO SCIENTIFIC PUBLISHING

INTERNATIONAL JOURNAL OF COMPUTER SCIENCE & INFORMATION SYSTEM
E-ISSN: 2536-7919
P-ISSN: 2536-7900

https://doi.org/10.55640/ijcsis/Volume11Issue08-06 PAGE NO: 59-67

Adaptive Graph Neural Network Model for Real-Time Intrusion and Cyber Threat Detection in Cloud Networks

Muhammad Ahmed Khan
School of Electrical Engineering
and Computer Science, NUST, Pakistan

Ayesha Noor
Research Department of Computer
Science, COMSATS University Islamabad, Pakistan

ARTICLE INFO

Article history:
Submission: June,27 2026
Accepted: July 18, 2026
Published: August 13, 2026
VOLUME: Vol.11 Issue 08 2026

Keywords:
Graph Neural Networks, Cloud Security, Intrusion Detection, Cyber Threat Detection, Adaptive Learning, Deep Learning, Network Security, Real-Time Analytics, Risk Analysis

ABSTRACT

The increasing interconnectivity, elasticity, and distributed architecture of cloud networks have created a security environment in which conventional intrusion detection approaches face difficulties in representing dynamic relationships among users, virtual machines, workloads, services, and network flows. This paper proposes an Adaptive Graph Neural Network (AGNN) Model for real-time intrusion and cyber threat detection in cloud networks. The proposed approach represents cloud infrastructure as a dynamic attributed graph and combines graph-based relational learning, adaptive representation updating, attention mechanisms, transfer-learning principles, and risk-oriented classification. The methodological foundation is informed by the data-centric perspective of machine learning (Emmert-Streib & Dehmer, 2022), deep transfer learning (Bashath et al., 2022), attention-based neural architectures (Cui et al., 2016), and transformer-based representation learning (Devlin et al., 2018; Clark et al., 2020). Experimental-design principles are incorporated to support systematic configuration and evaluation (Barker & Milivojevich, 2016; Cox & Reid, 2000). The framework further builds upon graph-based cyber-threat learning demonstrated for cloud platforms by Marri et al. (2025). The proposed model dynamically updates node and edge representations as network conditions change, enabling detection of anomalous communication patterns, compromised entities, lateral movement, and coordinated attacks. The findings indicate that a graph-oriented adaptive architecture provides a stronger theoretical basis for cloud intrusion detection than isolated feature-based classification because security decisions can incorporate both entity-level behavior and relational context. The principal contribution is an integrated framework that connects dynamic graph construction, adaptive graph representation, threat classification, and risk analysis within a real-time cloud-security pipeline.

INTRODUCTION

1.1 Background

Cloud computing has transformed information infrastructure from relatively static network environments into highly distributed ecosystems containing virtual machines, containers, APIs, storage resources, applications, users, service accounts, and dynamically established communication channels. Such environments produce security data characterized by high dimensionality, temporal variation, heterogeneous entities, and strong interdependencies. Consequently, intrusion detection cannot be considered exclusively as a problem of classifying individual network records. A security event affecting one entity may become meaningful only when its relationships with other entities are considered.

Traditional machine-learning approaches generally represent observations as independent feature vectors. This representation is useful when the predictive information is primarily contained within individual observations, but it becomes restrictive when the security meaning of an event depends on communication relationships. A compromised virtual machine communicating with several previously unrelated services, for example, can represent a substantially different threat from the same machine producing an isolated abnormal request. A graph representation can explicitly preserve such dependencies.

Graph-based deep learning therefore provides an appropriate theoretical foundation for cloud-security analytics. The work of Marri et al. (2025) demonstrates the relevance of graph-based deep learning for identifying cyber threats in cloud platforms. Building upon this direction, the present research introduces an adaptive perspective in which graph representations are not treated as permanently fixed structures. Instead, the proposed model continuously accommodates changing cloud relationships, behavioral patterns, and threat characteristics.

The conceptual direction is also consistent with the supplied Graph Deep Learning-Driven Cyber Threat Detection and Risk Analysis Framework for Secure Cloud Environments, which emphasizes the importance of combining graph-based learning with cyber-threat detection and risk analysis. In the proposed model, this principle is extended through adaptive graph updating and real-time threat classification.

1.2 Problem Statement

Cloud intrusion detection faces three interconnected challenges. First, cloud infrastructure is dynamic, meaning that nodes and communication edges can appear, disappear, or change characteristics. Second, cyber threats frequently exploit relationships among entities rather than isolated activities. Third, threat distributions may change over time, reducing the reliability of models trained under static assumptions.

Deep learning models can learn complex representations, but their effectiveness depends strongly on the quality and organization of data. Emmert-Streib and Dehmer (2022) emphasize the importance of a data-centric perspective in understanding machine-learning paradigms, while Bashath et al. (2022) demonstrate that transfer learning can be affected by characteristics of source and target data. These observations are particularly important in cloud security because different cloud deployments can exhibit substantially different traffic and behavioral distributions.

The research problem is therefore defined as follows: How can a graph-based deep-learning architecture dynamically represent changing cloud relationships and use these representations for real-time intrusion and cyber-threat detection while maintaining adaptive classification capability?

1.3 Research Objectives

The primary objective is to develop a conceptual Adaptive Graph Neural Network model capable of detecting and classifying cyber threats in dynamic cloud networks.

The specific objectives are to establish a graph representation of cloud entities and interactions; develop adaptive node and edge representations; incorporate attention-based relational learning; classify anomalous behavior into threat categories; generate risk-oriented outputs; and establish an evaluation methodology based on controlled experimental-design principles.

1.4 Scope and Significance

The research focuses on network and infrastructure-level cloud-security detection. The proposed framework can represent users, workloads, virtual machines, services, applications, and network flows as graph entities. It is intended for real-time or near-real-time analytical environments rather than offline forensic analysis alone.

Its significance lies in moving from isolated-event detection toward relationship-aware threat intelligence. The proposed framework does not claim that graph learning automatically resolves all cloud-security

challenges. Rather, it provides an architecture in which relational information, adaptive learning, and risk assessment can be integrated systematically.

LITERATURE REVIEW

Deep learning has developed through architectures designed to learn increasingly expressive representations from complex data. Devlin et al. (2018) demonstrated the effectiveness of bidirectional transformer representations for language understanding, while Clark et al. (2020) introduced ELECTRA's discriminator-oriented pre-training strategy. Although these studies focus on textual data, their broader methodological significance is the learning of contextual representations rather than relying exclusively on manually constructed features. For cybersecurity, contextual representation is similarly important because an individual network event may acquire meaning from surrounding activity.

Attention mechanisms provide another relevant theoretical foundation. Cui et al. (2016) demonstrated that attention-over-attention architectures can identify important interactions between information representations. This principle can be transferred conceptually to cloud-security graphs, where different neighboring nodes and communication edges should not necessarily contribute equally to a node's threat representation.

The data-centric perspective developed by Emmert-Streib and Dehmer (2022) is particularly relevant to graph-based security detection. Machine-learning performance depends not only on architecture but also on the properties, organization, and quality of data. In cloud networks, graph construction therefore becomes a fundamental modeling decision. Incorrectly defined nodes, missing interactions, noisy edges, or inconsistent temporal information can undermine downstream graph learning regardless of model sophistication.

Bashath et al. (2022) provide a broader analysis of deep transfer learning from a data-centric perspective. Their work highlights the importance of relationships between source and target data distributions. This issue is significant for cloud intrusion detection because a model trained in one cloud environment may encounter different workloads, communication structures, and attack distributions in another. An adaptive architecture should consequently accommodate distributional changes rather than assuming complete stationarity.

Back et al. (2020), Briggs (2021), and Chen et al. (2017) focus on question answering and machine reading comprehension. Their direct application is outside cybersecurity, but they contribute relevant insights into contextual representation, answerability, and information retrieval. In a cyber-threat environment, analogous principles can be used to distinguish informative contextual signals from insufficient evidence. The model should therefore avoid interpreting every anomaly as a confirmed intrusion.

Conneau and Lample (2019) demonstrate the value of cross-lingual representation learning for heterogeneous language environments. While the application domain differs, the underlying principle of learning transferable representations is relevant to heterogeneous cloud environments in which different workloads and services produce different behavioral distributions.

Barker and Milivojevich (2016) and Cox and Reid (2000) establish the theoretical importance of experimental design. For cybersecurity research, this implies that model comparisons should not depend on arbitrary configurations. Factors such as graph construction, learning rate, neighborhood depth, temporal window, feature selection, and class distribution should be controlled systematically.

The work by Marri et al. (2025) is the most directly aligned supplied reference because it explicitly addresses graph-based deep learning for cyber-threat identification in cloud platforms. It establishes a foundation for considering graph structures as useful representations of cloud-security relationships. The present research extends this direction by emphasizing adaptation, continuous graph evolution, and real-time threat interpretation.

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

Cloud Environments similarly supports the conceptual integration of graph learning with threat detection and risk analysis. However, a remaining gap is the explicit treatment of graph evolution as a first-class component of the detection process. The proposed AGNN model addresses this gap by allowing representations and threat assessments to respond to changing relationships.

METHODOLOGY

3.1 Proposed Research Architecture

The proposed architecture consists of six sequential functional layers:

Cloud Data Acquisition → Dynamic Graph Construction → Adaptive Graph Representation → Relational Attention Learning → Threat Classification → Risk Analysis and Alert Generation

The architecture is designed as a continuous pipeline rather than a one-time training process. Cloud telemetry enters the system through network flows, authentication events, service interactions, workload behavior, and infrastructure-level observations. These observations are converted into graph entities and relationships before being processed by the neural architecture.

The framework follows the graph-based cybersecurity direction established by Marri et al. (2025) while incorporating adaptive mechanisms intended for changing cloud conditions. The supplied Graph Deep Learning-Driven Cyber Threat Detection and Risk Analysis Framework for Secure Cloud Environments further motivates the integration of threat detection and risk analysis rather than treating classification as the final security objective.

3.2 Dynamic Cloud Graph Construction

Let the cloud environment at time t be represented as:

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

where V_t represents cloud entities, E_t represents relationships, and X_t represents node and edge attributes.

Nodes may represent users, virtual machines, containers, services, databases, applications, or cloud endpoints. Edges may represent network communication, authentication, resource access, API invocation, file transfer, or service dependency.

An edge can be represented as:

$$e_{ijt}=(v_i,v_j,a_{ijt})$$

where a_{ijt} contains relational characteristics such as communication frequency, protocol, volume, direction, and temporal information.

This formulation is important because cyber threats frequently manifest as changes in relationships. For example, a workload that normally communicates with three internal services but suddenly establishes connections with multiple unfamiliar endpoints creates a graph-level behavioral change that may be less visible in isolated traffic features.

3.3 Adaptive Node Representation

Each node initially receives an embedding:

$$h_i(0)=f(X_i)$$

where f is an initial feature transformation.

The adaptive component updates node representations as new observations arrive:

$$h_{i}^{t+1}=U(h_{i}^{t},m_{i}^{t},\Delta G_{t})$$

where m_i^t represents information aggregated from neighboring nodes and ΔG_t represents graph changes.

This mechanism enables the model to distinguish persistent behavior from newly emerging behavior. A sudden change in a node's neighborhood can therefore modify its representation without requiring the entire cloud graph to be reconstructed conceptually from scratch.

3.4 Relational Message Passing

For each node i, information is aggregated from its neighborhood N(i):

$$m_{i}=\sum_{j\in N(i)}\alpha_{ij}W h_{j}$$

where W is a learnable transformation and α_{ij} represents the importance assigned to neighboring node j.

An attention mechanism is appropriate because not all interactions have equal security significance. This theoretical approach is related to the attention principles demonstrated by Cui et al. (2016), in which learned attention weights help identify informative contextual relationships.

A multi-layer architecture can then be expressed as:

$$h_{i}^{(l+1)}=\sigma(W_{l}h_{i}^{(l)}+m_{i}^{(l)})$$

where σ is a nonlinear activation function.

3.5 Temporal Adaptation

A static graph may become inaccurate as cloud infrastructure changes. Therefore, the proposed model introduces temporal graph snapshots:

$$G_{t-k},\ldots,G_{t-1},G_{t}$$

The model compares recent structural states to identify behavioral deviations.

For example, repeated authentication failures involving the same account may initially appear as independent events. However, when these events are connected temporally and relationally with several services, they can represent a coordinated credential-abuse pattern.

This temporal design complements the data-centric perspective of Emmert-Streib and Dehmer (2022), because the usefulness of a feature is evaluated in relation to its evolving data context rather than assumed to remain constant.

3.6 Threat Classification

The final graph representation is passed to a classification function:

$$\hat{y}_{i}=\text{Softmax}(W_{c}h_{i}+b_{c})$$

where $\hat{y}_i$ represents the probability distribution over threat classes.

Possible analytical categories include benign activity, reconnaissance, credential-related attacks, denial-of-service behavior, lateral movement, malicious service interaction, and other anomalous activity defined according to the available training labels.

The model should distinguish between anomaly detection and confirmed classification. This distinction is

theoretically important because contextual machine-learning systems can encounter insufficient evidence. Insights from answerability-oriented research such as Back et al. (2020) provide a useful conceptual analogy: when evidence is inadequate, confidence should be treated differently from a definitive prediction.

3.7 Risk Scoring

Detection alone does not determine operational priority. The proposed framework therefore generates a risk score:

$$R_i = \lambda_1 P_i + \lambda_2 S_i + \lambda_3 C_i + \lambda_4 T_i$$

where P_i represents predicted threat probability, S_i represents severity, C_i represents contextual criticality, T_i represents temporal or behavioral escalation, and λ values represent configurable weights.

This enables security teams to prioritize events rather than responding identically to every detected anomaly. The risk-analysis orientation is consistent with the supplied Graph Deep Learning-Driven Cyber Threat Detection and Risk Analysis Framework for Secure Cloud Environments, while the adaptive formulation extends the concept to changing graph states.

3.8 Transfer and Adaptation Strategy

Because cloud environments differ, direct deployment of a model trained in one environment may lead to distributional degradation. Bashath et al. (2022) demonstrate why transfer learning must be examined through the characteristics of source and target data.

The proposed model therefore uses an adaptation stage in which previously learned graph representations are fine-tuned using new environmental data. Transformer-based representation learning from Devlin et al. (2018) and Clark et al. (2020), together with cross-domain representation principles from Conneau and Lample (2019), provides conceptual support for transferable representation learning.

3.9 Experimental Design

Evaluation should follow controlled experimental-design principles rather than relying on a single configuration. Barker and Milivojevich (2016) emphasize quality through experimental design, while Cox and Reid (2000) provide theoretical foundations for systematic experimentation.

Relevant factors include graph window size, number of graph-convolution layers, attention configuration, learning rate, embedding dimension, temporal update frequency, and adaptation strategy. Comparative experiments should evaluate static graph learning against adaptive graph learning to determine whether dynamic updating provides measurable benefits.

Performance should be assessed using precision, recall, F1-score, false-positive rate, detection latency, and computational overhead. Importantly, imbalanced threat classes should be analyzed separately because aggregate accuracy can obscure poor detection of rare but operationally important attacks.

RESULTS

The proposed methodology produces several analytically significant findings. First, graph representation provides a stronger conceptual representation of cloud-security behavior than isolated event modeling because it preserves relationships among infrastructure entities. A network connection, authentication event, or API invocation becomes more informative when evaluated in relation to neighboring activities. This finding aligns with the graph-oriented cyber-threat detection direction established by Marri et al. (2025).

Second, adaptive graph updating is theoretically expected to improve robustness in dynamic environments. A static model assumes that the relationships among cloud entities remain sufficiently stable, whereas real cloud infrastructures continuously change. The proposed temporal mechanism addresses this limitation by

incorporating new edges, removed edges, and altered interaction patterns into subsequent representations.

Third, attention-based aggregation provides a mechanism for distinguishing high-value security relationships from routine interactions. Rather than assigning identical importance to all neighboring nodes, the model learns relational weights. This is consistent with the broader theoretical utility of attention mechanisms identified by Cui et al. (2016).

Fourth, the framework indicates that detection and risk prioritization should be treated as related but distinct tasks. A classifier may identify an event as suspicious without determining whether it represents an immediate operational risk. The proposed risk score combines prediction confidence with contextual and temporal characteristics, thereby supporting more meaningful alert prioritization.

Fifth, data characteristics remain a fundamental limitation. Emmert-Streib and Dehmer (2022) and Bashath et al. (2022) indicate why data quality and distribution are central to machine-learning performance. In cloud security, incomplete telemetry, noisy graph edges, changing workloads, and insufficient examples of rare attacks can reduce generalization. Therefore, increasing model complexity alone cannot guarantee better detection.

Sixth, experimental evaluation must explicitly consider transferability. A model that performs well in one cloud environment may encounter different graph structures elsewhere. The transfer-learning perspective of Bashath et al. (2022) consequently supports evaluating adaptation rather than reporting only within-environment performance.

Finally, the proposed framework suggests that graph-based detection is most valuable when it is integrated into a broader security decision pipeline. The Graph Deep Learning-Driven Cyber Threat Detection and Risk Analysis Framework for Secure Cloud Environments reinforces this integrated perspective, while the present architecture adds explicit adaptive graph processing. Thus, the principal finding is not merely that GNNs can classify threats, but that adaptive relational representations can provide a more appropriate modeling basis for evolving cloud-security conditions.

Because the present study is a methodological framework rather than a report of a newly executed benchmark experiment, numerical performance values are intentionally not fabricated. Actual precision, recall, F1-score, latency, and computational measurements should be reported after implementation and evaluation on an appropriate cloud-security dataset.

DISCUSSION

The proposed AGNN architecture addresses an important conceptual limitation of conventional intrusion detection: the assumption that security events can be interpreted independently. In cloud environments, relationships among users, workloads, services, and endpoints are often central to identifying malicious behavior. A graph-based representation therefore provides a natural structure for capturing dependencies that conventional tabular representations may obscure.

The principal theoretical contribution is the integration of adaptation with graph representation learning. Marri et al. (2025) establish the relevance of graph-based deep learning for cloud cyber-threat identification, while the proposed model extends this foundation by treating graph evolution as part of the detection problem. This distinction is important because cloud infrastructure is not a static collection of entities. Its topology and behavioral patterns change continuously.

The attention component also has practical implications. Attention-based message passing can potentially reduce the influence of irrelevant neighbors while emphasizing relationships associated with anomalous behavior. This is conceptually consistent with Cui et al. (2016), although the present application involves graph relationships rather than reading-comprehension representations.

A second important implication concerns transferability. Bashath et al. (2022) emphasize data-centric considerations in deep transfer learning. Cloud-security models face a similar challenge because training

and deployment environments may differ substantially. Adaptive updating can reduce, but cannot eliminate, this domain-shift problem. Consequently, the proposed architecture should not be interpreted as universally transferable without additional validation.

The data-centric position of Emmert-Streib and Dehmer (2022) further suggests that graph quality deserves as much attention as neural architecture. Poor graph construction can introduce misleading relationships, while missing telemetry can create incomplete representations. Therefore, preprocessing, feature validation, entity resolution, and temporal alignment are critical methodological components rather than secondary engineering tasks.

Experimental validity is another concern. Barker and Milivojevich (2016) and Cox and Reid (2000) indicate the importance of controlled experimental design. For the proposed model, performance comparisons should isolate the contribution of adaptive graph learning from other factors. A fair study should compare static and adaptive configurations under equivalent training conditions and evaluate both detection quality and computational cost.

The framework also introduces a practical trade-off between detection sensitivity and alert volume. More aggressive anomaly detection may improve recall but increase false positives. Risk scoring can partially address this issue by prioritizing alerts according to contextual severity rather than simply generating binary alarms.

Nevertheless, several limitations remain. Real-time graph processing can introduce substantial computational requirements as the number of entities and interactions increases. Dynamic graph maintenance may also become challenging in highly elastic cloud environments. Furthermore, adversaries may deliberately manipulate relationships to evade graph-based detection. Model interpretability is another concern because security analysts may require explanations for why a particular graph neighborhood produced a high-risk prediction.

The Graph Deep Learning-Driven Cyber Threat Detection and Risk Analysis Framework for Secure Cloud Environments provides a useful conceptual basis for connecting graph learning with risk analysis; however, the present work emphasizes adaptive representation and temporal graph evolution as additional mechanisms for operational cloud environments. The resulting architecture should therefore be viewed as a research framework requiring empirical validation rather than as a universally optimal security solution.

CONCLUSION

This paper proposed an Adaptive Graph Neural Network Model for Real-Time Intrusion and Cyber Threat Detection in Cloud Networks. The framework models cloud infrastructure as a dynamic attributed graph and integrates adaptive graph representation, relational attention, temporal updating, threat classification, and risk scoring.

The principal contribution is the recognition that cloud-security detection should account for both individual behavioral characteristics and relationships among infrastructure entities. Graph-based learning provides the relational foundation, while adaptive updating addresses the changing nature of cloud environments. The approach is theoretically supported by research on graph-based cyber-threat learning, attention mechanisms, data-centric machine learning, transfer learning, and experimental design.

The model also identifies an important distinction between detecting suspicious activity and prioritizing security risk. Combining threat probability with contextual and temporal information can provide a more operationally meaningful security output than classification alone.

Future research should implement the framework using real-world or rigorously constructed cloud-security datasets, compare static and dynamic GNN architectures, investigate concept drift, evaluate computational scalability, and develop explainable mechanisms for analyst-oriented threat interpretation. Particular attention should also be given to cross-environment transfer and adversarial manipulation of graph structures. Such empirical validation would establish whether the theoretical advantages of adaptive

graph learning translate into measurable improvements in real-time cloud intrusion detection.

REFERENCES

1. Back, S., Chinthakindi, S. C., Kedia, A., Lee, H., & Choo, J. (2020). Neurquri: Neural question requirement inspector for answerability prediction in machine reading comprehension. In International Conference on Learning Representations.

2. Barker, T. B., & Milivojevich, A. (2016). Quality by experimental design. CRC Press.

3. Bashath, S., Perera, N., Tripathi, S., Manjang, K., Dehmer, M., & Emmert-Streib, F. (2022). A data-centric review of deep transfer learning with applications to text data. Information Sciences, 585, 498–528.

4. Briggs, J. (2021). Fine-tuning with squad 2.0.

5. Chen, D., Fisch, A., Weston, J., & Bordes, A. (2017). Reading wikipedia to answer open-domain questions. arXiv preprint arXiv:1704.00051.

6. Clark, K., Luong, M.-T., Le, Q. V., & Manning, C. D. (2020). Electra: Pre-training text encoders as discriminators rather than generators. arXiv preprint arXiv:2003.10555.

7. Conneau, A., & Lample, G. (2019). Cross-lingual language model pretraining. Advances in Neural Information Processing Systems, 32.

8. Cox, D. R., & Reid, N. (2000). The theory of the design of experiments. Chapman and Hall/CRC.

9. Cui, Y., Chen, Z., Wei, S., Wang, S., Liu, T., & Hu, G. (2016). Attention-over-attention neural networks for reading comprehension. arXiv preprint arXiv:1607.04423.

10. Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.

11. Emmert-Streib, F., & Dehmer, M. (2022). Taxonomy of machine learning paradigms: A data-centric perspective. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 12(5), e1470.

12. 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.

.