National College of Ireland

Policy-Driven Identity and Access Management: Strengthening Compliance with Security Standards and Emerging Technologies

MSc Research Project
MSc in Cybersecurity

Sreeram Krishna
Student ID: X23292431

School of Computing
National College of Ireland

Supervisor: Prof. Michael Pantridge

National College of Ireland

MSc Project Submission Sheet

School of Computing

Student Name:	Sreeram Krishna		
Student ID:	X23292431		
Programme:	MSc in Cybersecurity	**Year:**	2024-2025
Module:	MSc Research Project		
Supervisor:	Prof. Michael Pantridge		
Submission Due Date:	11-08-2025		
Project Title:	Policy-Driven Identity and Access Management: Strengthening Compliance with Security Standards and Emerging Technologies		
Word Count:	7129	**Page Count:**	23

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.
ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature:	Sreeram Krishna
Date:	11-08-2025

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST

Attach a completed copy of this sheet to each project (including multiple copies)	□
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	□
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	□

Assignments that are submitted to the Programme Coordinator Office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Policy-Driven Identity and Access Management: Strengthening Compliance with Security Standards and Emerging Technologies

Sreeram Krishna
X23292431

Abstract

As more organizations choose the cloud, secure user access is vital for stopping breaches, unauthorized access, and account hijacking. Most IAM systems of this kind are not dynamic and depend on roles, which makes it hard to detect all forms of suspicious behavior. This study presents the Graph Neural Network with Temporal Attention (GNN-TA) model for noticing unusual behavior of users in cloud-based Identity and Access Management systems. Data from the CLUE-LDS, which covers real user activity for five years, are used by the model to form a graph that reflects user behavior and gives more importance to unusual activity over the recent period. Synthetic examples of hijacking and similar risks are added to the data to test the model's ability to notice such troublesome actions. The approach makes it possible for IAM to spot threats as they emerge by combining adaptation and policy-based actions. The findings are used to develop cloud security systems that are aware of privacy needs and match the latest security regulations.
Keywords- Mobile device security, Malware, Data theft, MAM, Endpoint security, SIEM

1 Introduction

With the rapid development of digital infrastructure, which is provoked by cloud computing, technologies of remote access, and hybrid environments, the general principles of user identities management and access control to resources in organizations have shifted (Dawood *et al.*, 2023). The importance of Identity and Access Management (IAM) as a pillar of cybersecurity has increased in recent times, as it takes care of the fact that only authorized people can receive the correct data, systems, and services. Conventionally, IAM systems have been implemented using a role-based or rule-based model, acquisition of access routes using premature permits, and occasional auditing to handle access (Singh, Thakkar and Warraich, 2023). Although such strategies have been satisfactory in a relatively controlled experience on-premises, they are vulnerable and encounter big difficulties in the contemporary dynamic, decentralized, and high-volume environment. New complexities and additional expansion of the attack surface have been caused by cloud platforms, multi-device access, and increased integrations with third parties (Azad *et al.*, 2024). This has made the historical models of IAM susceptible to losing ground on the intricacies that matter in contextual and behavioral terms to identify advanced cyberattacks like insider attacks, account takes, and policy noncompliance (Felix Chad, 2025). Security policies are supposed to offer a governance layer to govern such risks nevertheless when there are no smart enforcement measures, security policies can be made ineffective. To address these issues, artificial intelligence (AI) and machine learning (ML) have become popular in IAM, which include remedies that are flexible and data-driven rather than

rule-based (Adenola, 2023). Within such approaches, graph-based anomaly detection techniques, especially Graph Neural Networks (GNNs), are emerging to be incredibly powerful in capturing the intricate graphs of relationships between users, devices, actions and access timelines. In combination with temporal attention models, it is possible to apply them to capture behavioral patterns over long periods and detect latent signs of emergent threats. The combined policy-based IAM and built on AI-driven analytics can create a new course as organizations are under growing regulatory and operational pressure to secure their systems. This paper is to trace that intersection, offering a model that enhances not only the access control but also the threat detection in contemporary IAM situations.

1.1 Importance of the Study

In the modern cybersecurity world, Identity and Access Management (IAM) has become mission-critical in the context of protecting organizational assets, data privacy, secure digital transformation. Due to the increased use of cloud volumes, distributed working environments, and interdigitated online networks, the technology has become a primary security need to understand who has access to what, when, and on what circumstances (Siraparapu and Azad, 2024). This makes IAM important in dealing with insider risks and keeping unauthorized access at bay, and ensuring Integrity by enforcing authentication, authorization, and policy-based access control. Moreover, data protection regulations like GDPR, HIPAA, or ISO/IEC 27001 strictly require access control, auditability, and responsibility, and, therefore, IAM is a necessity to avoid being non-compliant. Nonetheless, traditional IAM solutions tend to be inflexible when supporting the identification of behavioral abnormalities and reacting to emerging threat vectors (Adebola Folorunso *et al.*, 2024). This highlights the increasing demand to have the context-sensitive and smart IAM solutions which could combine not only machine-learning but also policy-driven decision-making to enhance access-governance and be in compliance with the security and regulatory goals.

1.2 Research Objectives

1. To model IAM user access patterns in cloud environments using graph-based representations.
2. To design and implement a GNN-TA model that integrates behavioral context and temporal dynamics for anomaly detection.
3. To evaluate the GNN-TA model's performance against standard ML baselines using real and synthetic behavioral data.
4. To propose enhancements to IAM policies based on the behavioral insights uncovered through anomaly detection.

1.3 Research Questions

1. How can user behavior and access relationships in cloud IAM systems be effectively modeled as a dynamic graph?
2. To what extent does the integration of temporal attention improve the performance of anomaly detection compared to traditional models?
3. What types of anomalous behaviors are best detected using the GNN-TA framework?
4. How can behavioral anomaly insights be used to inform and optimize IAM policy compliance?

1.4 Research Problem Statement

Conventional forms of Identity and Access Management (IAM) systems, as useful as they are in implementing fixed access controls and authentication methods, have become ineffective in identifying dynamic, contextually aware anomalies, which emerge in multi-dimensional and dynamic cyber settings (Godwin Nzeako and Rahman Akorede Shittu, 2024). In majority of the conventional IAM systems, roles, permissions and rule-based policies have been the major components, which are not very flexible to the shift in user behaviors and access patterns. As companies move to cloud and consider remote or hybrid work arrangements, accessing information resources via a variety of devices, geo-locations, and time zones becomes a common phenomenon due to which behavioral baselines introduce new norms that cannot be monitored well by using static IAM systems (Himeur *et al.*, 2021). Typical IAM solutions, along with MFA and federated protocols, have difficulty detecting minor behavior changes that are associated with insider threats, hijacking or policy breaches. Such anomalies can be in the form of subtle changes in access patterns, not necessarily breach patterns. Consequently, existing solutions do not offer any proactive context and real-time threat detection mechanisms, which re necessitate intelligent and context-sensitive behavioral-based IAM solutions.

1.5 Scope and Limitations

The proposed study aims to optimize Identity and Access Management (IAM) by leveraging policy-driven frameworks and using more powerful anomaly detection constructions based on Graph Neural Networks with Temporal Attention (GNN TA). The study narrows to focus specifically on cloud based IAM systems, with structured log files like the CLUE-LDS data set which records other user behaviors like log in and file access as well as administrative tasks. The selection of the dataset indicates a controllable representation of cloud environment, which will be called realistic in order to come up with and test the proposed model of anomaly detection. Nevertheless, there are a number of limitations that have to be taken into consideration. First, it is based on account hijacking and insider threat simulated scenarios in order to assess model performance. Though such simulations are intended to reflect true attack trends in the real world, they fail to reflect fully on the nature of challenges that exist in enterprise operation due to the more complexity, unpredictability and plurality. Second, the proposed model implementation and testing is piece meshed in a sand box experiment environment as opposed to a live production IAM environment. This can also restrict the cross-applying of results to big-sized or quickly changing enterprise situations. Regardless of these limitations, the research provides a beneficial proof of concept framework that can guide the future development of the intelligent, policy-aware IAM systems.

1.6 Research Outline

Provide a brief overview of each chapter:

- Chapter 2: Literature review on IAM, anomaly detection, policy models, and GNN applications.
- Chapter 3: Methodology covering data source, model design, and evaluation criteria.
- Chapter 4: System implementation and architecture description.
- Chapter 5: Results and evaluation of the anomaly detection system.
- Chapter 6: Discussion of findings and implications.
- Chapter 7: Conclusions and future research directions.

2 Literature Review

Identities and access controls that are either digital or incorporated into digital access are handled by modern systems of IAM through authentication (e.g., passwords, biometrics, MFA) and authorization (e.g., RBAC, ABAC, PBAC). As cloud, mobile workforces, and third-party integrations have become a reality, IAM has been pushed beyond on-premise boundaries in support of hybrid environments. It is no longer limited to specific user populations: employees, partners, customers, and maintains both security, compliance, and user experiences. The newer paradigms such as IDaaS and Zero Trust are introducing plans of smarter, context-aware IAM, but most of the systems continue using more strict policies.

As identity theft increases, it seems to be a logical step to examine how AI can improve fraud detection capabilities within IAM systems by detecting complex patterns in the large, dynamic data. (Tamraparani, 2023) emphasizes such advantages as enhanced accuracy, and also there are disadvantages such as false alarms, inconsistency in data, and compatibility to existing systems. As cyber threats increase, the combination of User Behavior Analytics (UBA) and Identity & Access Management (IAM) becomes helpful in real-time threat identification and control of access. In this paper, the focus is on the ability of AI, machine learning, and blockchain to empower UBA-IAM solution against insider threats, APTs, and credential stealers. In spite of their potential, issues of false positives and scalability in large environments have been identified as areas of concern to development in the future (Vitla, 2023).

(Cheruku and Narasimha, 2024) proposes a one-dimensional Separable Convolutional Neural Network (1D-SCNN) to identify a fast and efficient way of detecting anomalies in Identity and Access Management (IAM) systems. Through the use of deep learning, the model is able to analyze user behavior and access patterns as well as examining methods, such as separable convolutions, separable convolutions, Leaky ReLU/ELU, MaxPooling, Dropout, and Flatten as a low computation overhead. 1D-SCNN structure is highly accurate (96%), as it uses past data to teach it how to recognise abnormal access behaviour. Nonetheless, it depends too much on historic trends, which are not effective in fending off new, sophisticated attacks, and also the reduced complexity of design might decrease the accuracy of dynamic IAM. The present study, (Folino, Otranto Godano and Pisani, 2023) presents a framework of real-time log analysis with ELK stack (Elasticsearch, Logstash, Kibana) and Kubernetes that will be used to track the behavior of users and anomalous events. It operates on snarling volumes of information on the internet and user dumps and categorises users depending on history of their digital marks of assorted data sources. The system is effective in detecting unusual patterns, solving incomplete data, and false alarm minimization. Its use, however, requires a high level of technical skills, especially the setting up and management of ELK and Kubernetes infrastructure. (Demirsoy *et al.*, 2024) proposes an AI based IAM that has a hybrid CNN-LSTM model to detect anomalies regarding user behaviors in real time. The input was gathered using a .NET worker service and pre-processed using user normalization obtaining precision and accuracies of 85.44, 87.95 and AUC of 0.8578. The model can provide scalable and adaptive IAM solutions yet issues exists in the concept of interpretability as well as in the process of enterprise data with a high dimensionality. (Sun, Yang and Zhang, 2020) proposed anomaly detection framework designed as a user centric framework is meant to detect a real time abnormal behavior of a user in secure Management Information Systems (MIS). By applying the web logs and clickstreams, multiple and mingled behavioral models are constructed, in which user behavior is assigned labels of roles and patterns in construction of separate and fused behavioral models. Through these models, real time searching and matching

of patterns and alerting is possible to determine any suspect activities. Although the system is effective when dealing with real time reaction, it does not perform well when dealing with such dynamic complex behaviors that are driven by outside factors**.**

Table 1: Summary of Literature Review

Ref	Problem Identified	Proposed Solution	Limitation
(Fan *et al.*, 2024)	Traditional GNNs (e.g., GCN, GAT) in IIoT systems fail to capture both spatial and temporal features.	STGaAN: Spatio-Temporal Gated Attention Networks combining gated GAT + temporal convolution + joint optimization via reconstruction and prediction loss.	GDN struggles with temporal feature learning; GAT lacks generalization; does not address data imbalance.
(Peng *et al.*, 2022)	BGP anomaly detection often ignores precise time dependencies and assumes static feature relationships.	Multi-view model using STL decomposition for noise reduction and GAT for feature and temporal correlations.	Relies on predefined relationships; may miss novel anomalies; effectiveness sensitive to real-world BGP variations.
(Yan, Luo and Shao, 2023)	Existing log-based methods struggle with time-aware relationships and graph structure for discrete events.	Temporal event graph + GCN with contrastive learning + TCN + GAN for future event link prediction.	Performance degrades with noisy/incomplete logs; complex model has high computational cost.
(Liu *et al.*, 2024)	Anomaly detection in distributed systems is hindered by complex temporal-logical dependencies.	TLAN: Temporal Logical Attention Network with multi-scale extraction and adaptive thresholds.	May face challenges in extremely large-scale dynamic environments with noisy log data.
(Guan *et al.*, 2022)	Multivariate time-series models often ignore temporal and inter-sensor correlation in smart factories.	GTAD: Combines TCN for temporal modeling and GAT for spatial sensor correlation.	Cannot explain the root cause of anomaly; interpretability is limited.
(Natha *et al.*, 2025)	Manual monitoring of surveillance video is inefficient and prior models fail to balance spatial-temporal anomaly detection.	CRBA: Combines DenseNet201 + BiLSTM + bi-attention to detect spatio-temporal anomalies.	High computational resource requirement; performance may drop in bad weather or limited-resource setups.
(Zhao *et al.*, 2024)	Performance degradation of anomaly detection models in high-dimensional multivariate time series.	Combination of Graph Attention Network (GAT) and Informer for short- and long-term forecasting anomaly detection.	Needs further optimization for richer time-series features; real-world ICS deployment not yet validated.
(Y. Zhang *et al.*, 2024)	Vessel anomaly detection methods often ignore temporal dependencies and feature correlations.	VBAD-GAT framework with time and feature graph attention modules and joint detection using reconstruction and prediction.	Dependent on training data quality and quantity; generalization to unseen scenarios may be limited.

(Z. Zhang *et al.*, 2024)	Lack of labeled data and difficulty in modeling system topology in cloud environments.	GCAD model combining GraphGRU, contrastive learning, and data augmentation for spatiotemporal anomaly detection.	Absence of annotated data hinders validation; capturing accurate system topology remains challenging.
(Ge *et al.*, 2024)	Existing models overly optimize time-series prediction, risking loss of critical anomaly information.	Spatiotemporal model with enhanced attention and heterogeneous graph contrastive learning to retain anomaly cues.	Effectiveness is highly dependent on quality annotated datasets; underrepresented anomalies may be mislearned.

The analysed articles reveal the impressive developments in the domain of anomaly detection, especially in systems that have a high degree of dimensionality, time complexity and spatial dependencies, namely IIoT networks, distributed systems, cloud infrastructures, and surveillance settings. The point of insufficiency of conventional models, such as GCN, GAT, or LSTM in capturing both temporal and spatial dynamics within multivariate temporal sequences, emerges as an outstanding theme throughout the literature. To counter this situation, researchers have proposed hybrid models to combine attention, graph neural networks (GNN) and temporal convolutional networks. As an example, STGaAN and GTAD combine graph attention and time-based modeling to identify anomalies in IIoT and smart factory environments, and VBAD-GAT and GCAD add more dedicated adjustments to maritime applications and cloud computing infrastructure, respectively, with the prevalence of contrastive learning and graph spatiotemporal reasoning. Though all these developments have been made, the limitations of the available models put the existence of a research gap at the fore. The present models usually require enormous labeled data, are cumbersome and their output in not interpretable thus not suitable to be used in real time. The dynamic threats cannot be adjusted using static models, and it is important to propose a simple system of anomaly detection which is flexible, lightweight, and explainable.

2.1 Research Gap

Identity and Access Management (IAM) systems are the first line of defence to implement access control and secure sensitive resources in the fast-changing, ever-evolving cloud environments today. However, with the expansion of the utilization of cloud services, the normal IAM-based protections, which commonly rely on printed rules and a set of pre-determined approaches to access, may fail to identify more advanced threats like insider abuse or takeover of an account (Ike *et al.*, 2025). The threats often occur in minute differences in user patterns, like unexpected access of files at unusual time and place. Regrettably, the majority of the available anomaly detection structures do not take into consideration the complex relationship between user roles, resource types, action patterns and time by being unable to consider them. They are either self-isolated time-series types or employ flat behavioral vectors, which creates a high false-positive environment coupled with a void in actionable intelligence (Himeur *et al.*, 2021). The central issue to be discussed in the current research emerges in the absence of a dynamic, context-sensible anomaly detection framework in the IAM systems able to distinguish between the legitimate variations of behavior and the real security threat. The current models fail to be able to model relational dependencies between objects (users, resources, services), and fail to capture the time changing profile of behavior

(Huang *et al.*, 2025). This renders them unready to be realized under real time detection on contemporary cloud ecosystems. The technical limitations that are encountered are the need to convert heterogeneous IAM log data into structured graph formats, the dynamic behavior variations, the ability to formulate and score anomalies and the need to integrate the anomaly detection process with policy-based enforcement mechanisms. Moreover, the critical constraints in terms of enterprise deployment are scalability and interpretability. To cover these gaps, the proposed research paper presents a new machine learning model Graph Neural Network with Temporal Attention (GNN TA) that models multi-entity user interaction over time and learns to emphasize suspicious variations according to structure and sequence. Through this framework, the project will set out to amplify how IAM systems can react to complicated behavioral anomalies in a prompt, specific, and explicable way, which can end up in a better secured and dynamic cloud security stance.

3 Research Methodology

The section introduces the research methodology used to create an effective anomaly detection system to identify anomalous activities in cloud-based Identity and Access Management (IAM). Its design is based on the opportunity of user behavior analytics by means of a Graph Neural Network with Temporal Attention (GNN-TA), capturing the patterns of the relationships as well as temporal changes in user activity. This approach implies data processing based on CLUE-LDS, the creation of a dynamic graph of user-resource pairs, and the extraction of meaningful behavioral characteristics, as well as the model training to detect deviations that can lead to security risks. Such a systematic methodology enables the proper and context-sensitive detection of abnormal patterns of accesses to the contemporary cloud settings.

3.1 Data Source

The data used in the present research is the CLUE-LDS (Cloud-based User Entity Behavior Analytics Log Data Set), a realistic and extensive background to study how users behave in the cloud setting. The dataset is obtained in more than five years (7 July 2017 to September 29 2022) and consists of about 50 million of log events created by more than 5,000 different users. The platform that these users were working on and interacting with was a real cloud storage system called hBOX, on which the Huemer Group-an IT service provider, based in Vienna, Austria-has been working and developing since early 2019 and extends the open-source framework Nextcloud with features like uploading of files, synchronization between devices, versioning, collaborative document management, and sharing. The manner in which users log on, retrieve files, share links, make configuration adjustments, and perform searches is highly represented in the log data and, therefore, offers a rich source of behavioral context to User and Entity Behavior Analytics (UEBA). Though the dataset largely represents regular user behavior, as its main advantage, it is especially applicable to the study on anomaly detection due to the ability to simulate the abnormal behavior, i.e., account hacking. In the quest to champion this, the repository of the dataset has seen the addition of utilities in analyzing similarities between users and behavior-switching tools which enable the researchers to emulate anomalous events by swapping the behaviour of one user with that of another. Although a few data were lost (109 missing events on three particular days because of database failures), the dataset was significantly complete and provided a rather firm basis of analysis and development of advanced anomaly detection models, including the one suggested in this paper, namely, GNN with Temporal Attention (*NIAID Data Discovery Portal*, 2022).

3.2 Preprocessing

In anticipation of a successful application of the CLUE-LDS dataset on anomaly detection modelling, a detailed data preprocessing procedure was applied to guarantee consistency, reliability and appropriateness of data towards behaviour analysis. Theoretically the first processes included deleting those files in the raw logs that were spoilt or incomplete, or matching, and aligning the time stamp to enable proper timing. Some activities like logging in, accessing files, sharing links and making configurations were identified and standardized. Based on this cleaned data, a series of crucial behavioral characteristics were determined such as logged in frequencies by user, time-of-day access probabilities, IP/geolocation variability, access to resource diversities, and file-sharing patterns. The aggregation period was used differently to represent short and long- term dynamics of behavior through the aggregation of these features. As the original dataset only consisted of benign activity, synthetic anomalies were injected to represent an attack on that data in a real-world situation, e.g. account hijacking or insider misuse. It was achieved by means of the given user-switching mechanism that substitutes the pattern of behavior of one user by another one and thus causes sudden jumps in activity that demonstrate instances of malicious access. These artificially injected anomalies made it possible to ensure that the data set could be used to critically evaluate the prescription of a proposed anomaly detecting model but maintain a realistic structure and context with actual production cloud IAM environment.

3.3 Model Development

The fundamental goal of the model development stage is to construct an intelligent anomaly detection system that will recognize suspicious user behavior in cloud-based Identity and Access Management (IAM) systems. This is done by a Graph Neural Network with Temporal Attention (GNN-TA), that not only utilizes the structural benefits of graph-based modeling, but also has the capability of predicting time-sensitive behavioral trends. This has been formulated by encoding the users and cloud resources as nodes and their interactions as time-stamped edges: this model learns complex relational and temporal dependencies in user activity. Through this strategy, the system can capture low-level anomalies that the traditional systems and rule-based or static systems are likely to ignore, including unauthorized access and policy breaches.

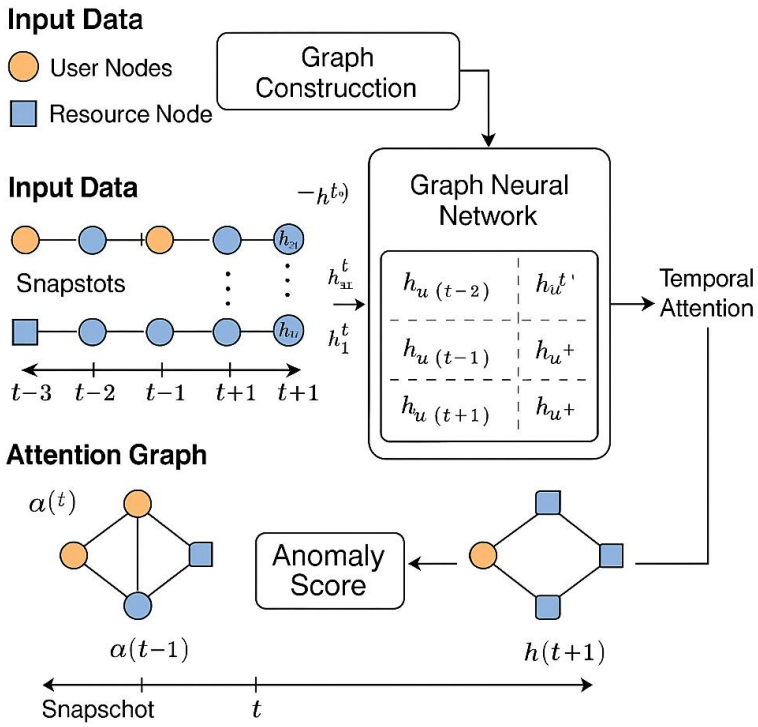

Figure 1. Workflow of GNN-TA Model

Graph Construction:

- Each **node** represents either a user or a cloud resource (e.g., files, configurations).
- **Edges** indicate access events (e.g., logins, file shares) between users and resources, with timestamp attributes.
- **Node features** include login frequency, time-of-day access patterns, location/IP variance, and access sequence history.

GNN TA Architecture:

- The GNN component learns spatial dependencies, such as which users regularly interact with which resources.
- A temporal attention layer weighs access events based on their recency and contextual significance.
- This enables the model to detect abnormal changes in behavior over time, such as sudden access to sensitive data at odd hours or from unusual locations.

Graph Node Attribute Design:

- Behavioral metrics (login time distributions, session duration, resource access patterns) are embedded into nodes.
- Temporal continuity is maintained to track progression and deviation in user behavior.

Graph Learning:

- The model iteratively updates node embeddings by aggregating information from neighbors and weighting time-sensitive interactions.
- Final embeddings are analyzed for deviations from learned behavioral norms to flag anomalies.

This combination of graphs and time allows a contextual and time-sensitive analysis of access behavior, so GNN- TA is more appropriate to detect possible security breaches in IAM systems.

3.4 Evaluation Metrics

Model performance is assessed using a combination of standard classification and anomaly detection metrics. These include:

- Accuracy – Proportion of correct predictions (normal vs. anomalous).
- Precision – Proportion of predicted anomalies that are true anomalies.
- Recall (Sensitivity) – Proportion of actual anomalies correctly identified.
- F1-Score – Harmonic mean of precision and recall.
- AUC-ROC – Area under the curve for true positive vs. false positive rate.

Moreover, the effects of identified anomalies on IAM policy enforcement are assessed. This entails identifying the extent to which the model can be used to assign the suspicious behavior to particular policy violation, say, an attempted access to the system without prior authorization or misuse of role. This situational analysis proves the pragmatic advantage of the model in enhancing access control implementation and consequent minimization of false positive in the IAM alerting schemes.

4. Design and Implementation Specifications

The section of the System Design and Implementation provides an architectural and technological basis of the proposed anomaly detection system, which will be designed to work in a contemporary cloud environment. The system is designed with the utilization of the Microsoft Azure cloud as the first layer of the infrastructure and drives scalability, reliability, and security of the data processing. The front-end interface is using Next. js which will present a responsive and interactive port of calling which will allow the administrators to visualize the behavioral anomalies and accordingly manage the IAM policies. The backend processing with behaviour analytics is based on Python and applies state-of-the-art graph neural networks, which process and analyze log records of user activity actions in the system with dedicated temporal attention. The application is cloud-native, enabling it to support cloud IAMs and cloud ecosystems as well as behavioral anomaly detection in real-time in line with increasing demands of the enterprise cloud adoption and secure digital operations.

Step 1: Microsoft Azure AD Authentication

This screen reflects the secure log in front end to connect to the Cloud IAM Dashboard at Microsoft Azure Active Directory (AD). It becomes a two-factor authentication (2FA) gateway, so only the person who has all the rights accesses the system. The user is required to log in using verified Microsoft credentials according to the level of identity management followed by the enterprise.

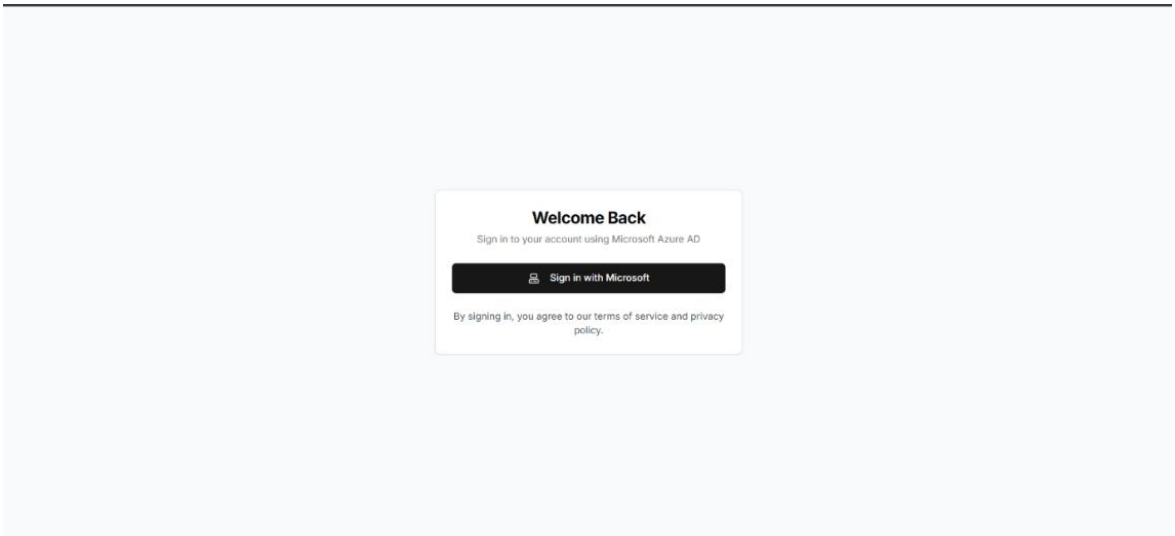

Step 2: IAM Dashboard Overview

After successful login, the user is directed to the Cloud IAM Dashboard. This dashboard shows a comprehensive overview of account activity, including:

- **User Profile Info** (name, email, role as administrator)
- **Security Assessment** powered by AI, highlighting:
    - **Risk Level:** Low Risk
    - **Risk Probability:** 10.0%
    - **Alert Message:** Your activity appears normal
- **Cloud Usage Metrics:** CPU, memory, and storage usage across subscriptions
- **Azure AD Group** membership for role-based access control

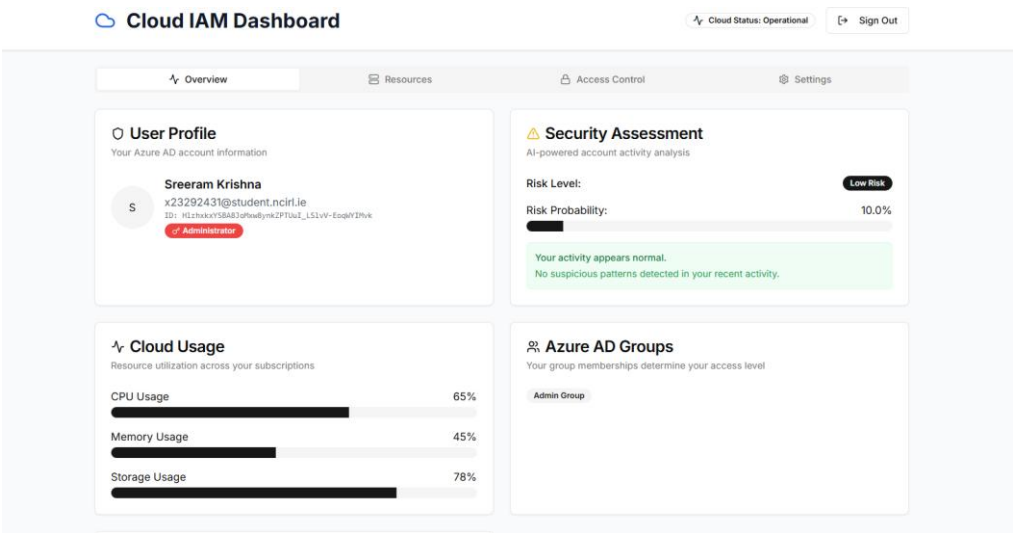

Step 3: Admin Panel

The admin panel provides core functionalities to manage the IAM system:

- Overview of **total users**, **active users**, and **system health**
- No unresolved security alerts
- Options for:
    - **User Management** (Manage users, permissions, and roles)
    - **Security Center** (Audit logs, security dashboard)
    - **Activity Monitoring** (Track user activity, generate reports)

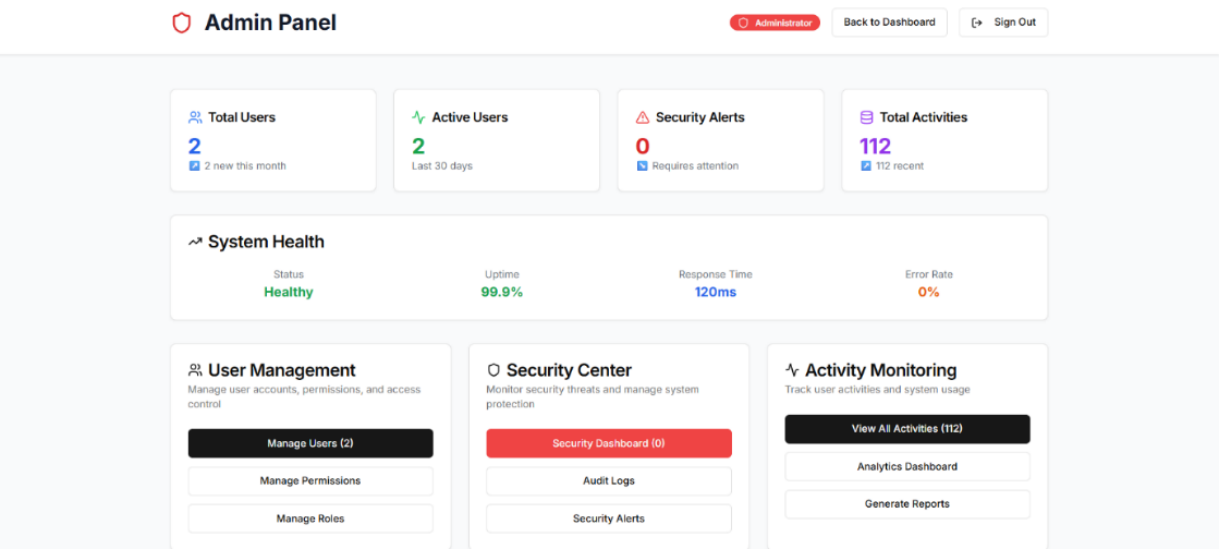

Step 4: User Management

This screen lists all registered users with metadata including:

- **Email and ID**
- **Status (Active)**
- **Activity logs**
- **Account age**

Admins can **view or delete** user accounts, enabling effective lifecycle management.

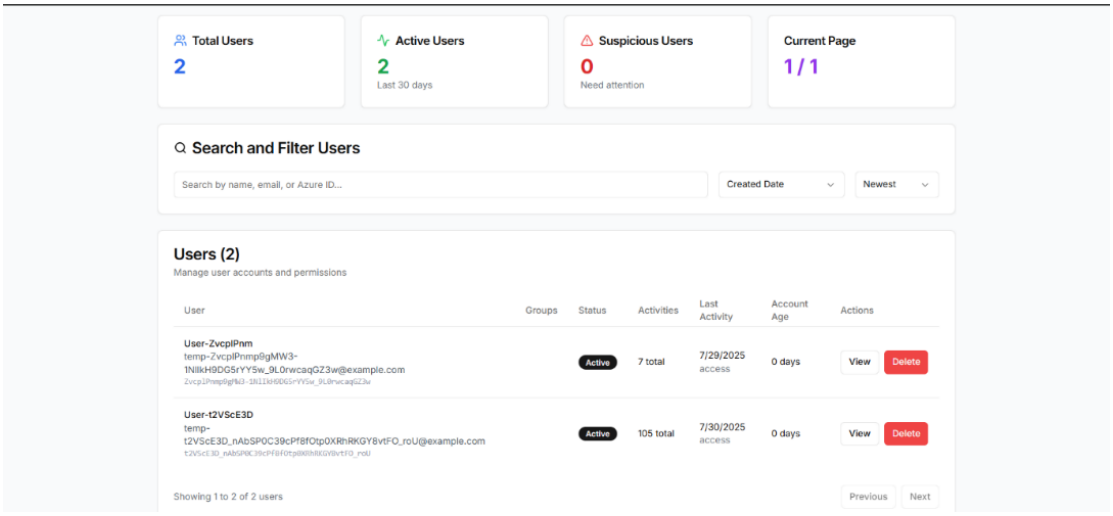

Step 5: User Profile and Quick Actions

This page shows full user details, with tools for:

- Enabling/disabling account status
- **Force password reset** for compromised accounts
- Tabs for user profile, activity, permissions, and security status

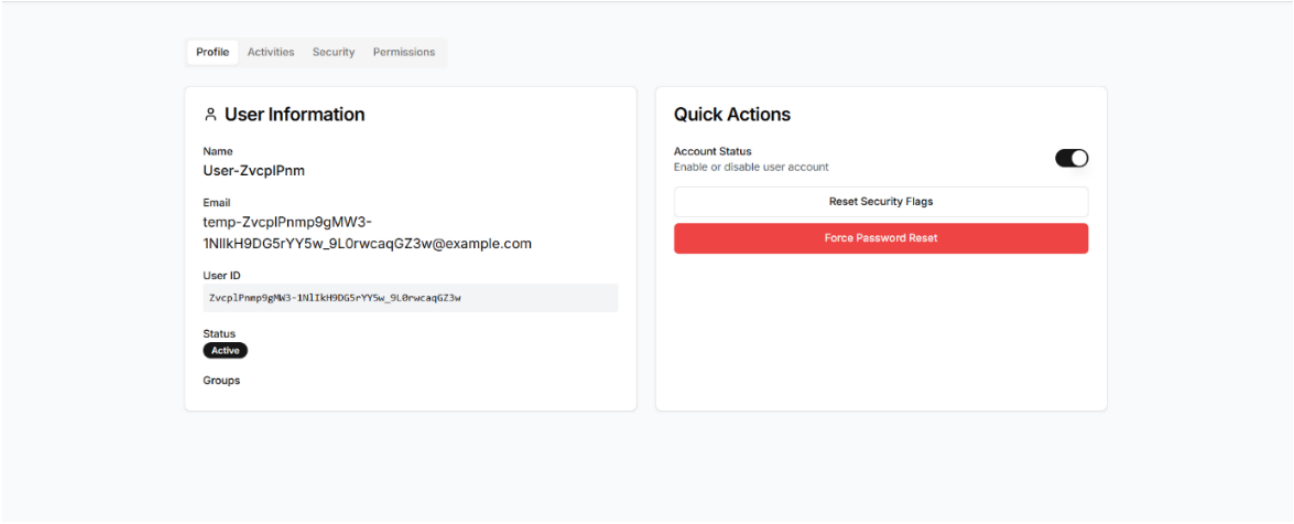

Step 6: Security Tab for Selected User

Displays detailed **security status for a selected user**:

- Indicates **no security alerts**
- Admin options include:
    - Viewing **login history**
    - Checking **failed attempts**
    - **Suspending the account** if necessary

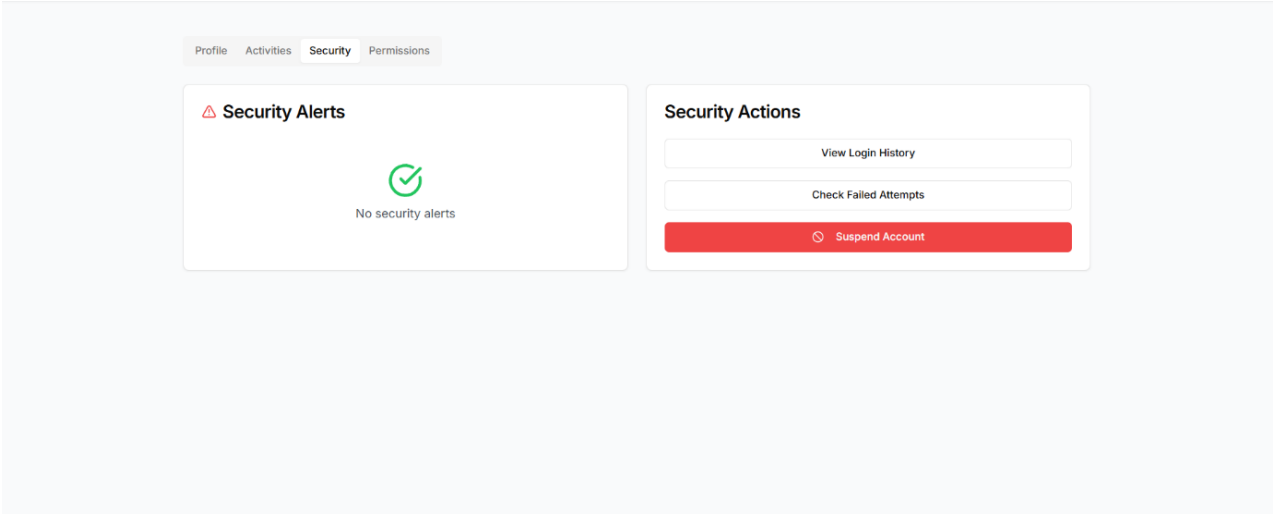

5. Implementation of Behavior-Aware Anomaly Detection

The figure shows the frequency of ten different system related user actions documented in the data set. By far the most frequent event is file_accessed, with almost 200,000 events, which means that users are interacting with the existing data on a high level. ill deposit1-file created occurs 174,954 times, and ill deposit1-file written 175,132 times, which shows the regular content creation and update. Relative scarcity can also be seen in other actions such as file_deleted or deleted_from_trashbin with corresponding counts of less than 30 000 and 25 000 respectively, apparently indicating that deletion and archival of data are subject to rare patterns. The popularity of the distribution focuses on a use pattern that is strongly biased in relation to viewing and editing data but not deleting.

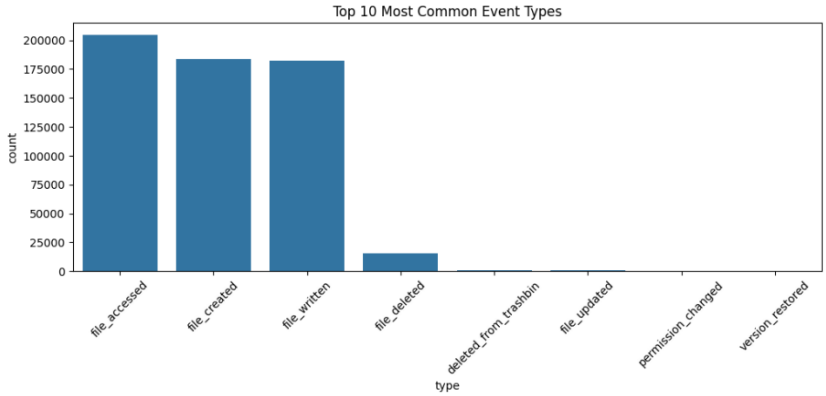

Figure 2. Distribution of User Activity Event Types in the Dataset

The bar graph displays the different levels of vulnerability to account hijacking to five user personas that include the external, management, technical, administration and sales. The hijack rate is highest on external users reaching to 0.200, which means that there is a high probability of hijack on the external users as compared to the internal positions. Moderate vulnerability rates were borne by management and technical users with 0.025 and 0.010 respectively whereas administration and sales roles remained unaffected with a rate of less than 0.005. Such distribution proves that much stronger security measures of external user accounts are necessary as it is clear that they are more likely to become sources of hijacking.

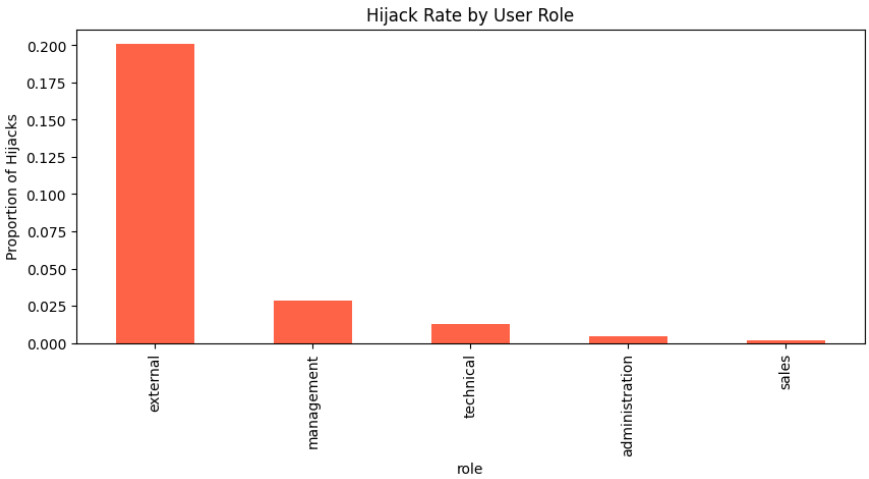

Figure 3. Hijack Rate by User Role

The horizontal bar chart emphasizes the top ten causative sources of attacked system resources during event hijack cases with each of them having three events respectively. The equal numbers among such resources as awful-crimson-dormouse-ga and immense-tan-ape-travelcon can lead to the conclusion that hijacking attempts are distributed instead of targeted on a chosen asset. This even highlights the need to have a strong system-wide security system since no discriminating attitude in terms of the type of resources seems to be evident on the attackers.

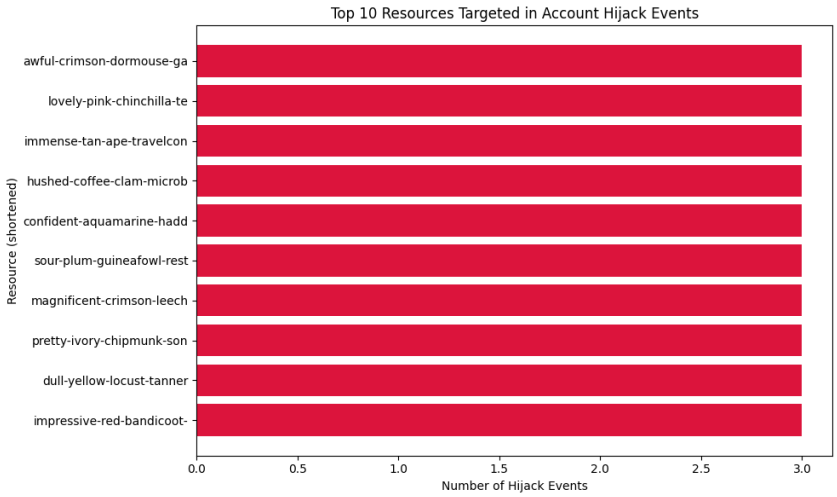

Figure 4. Top 10 Resources Targeted in Account Hijack Events

The network graph is a visual representation of the interactions within the system in the form of nodes that represent the users in the system and the edges represent types of connections between them: red color represents hijack events, and blue connectivity represents normal activity. It is worth noting that node 1 is only connected using red edges, which is an indication of it being a circumstantial hijack target whereas node 0 has a combination of both red and blue edges, therefore partially compromised or exposure. This comparison can give the user account of patterns susceptibility and will help to locate possible source of threats on the network.

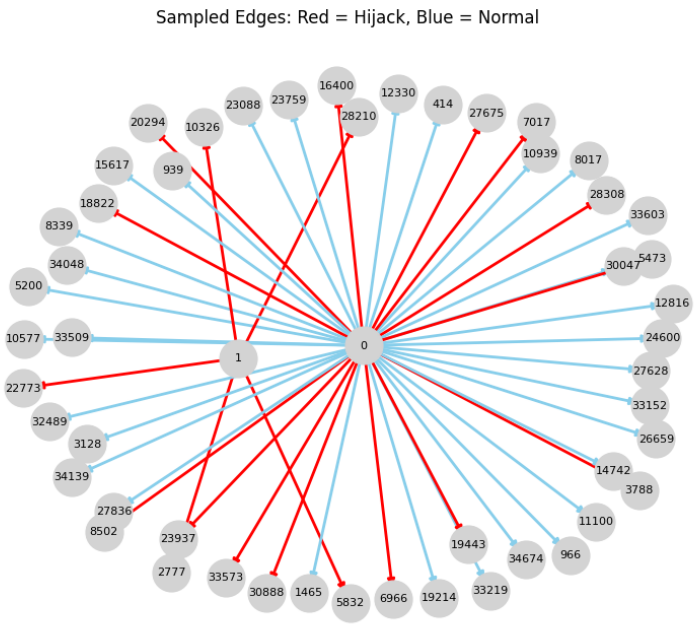

Figure 5. Sampled Edges – Red = Hijack, Blue = Normal

Figure 6 monitor the performance of the model with 25 training epochs. Accuracy over Epochs - As can be seen, the training accuracy approaches 1.0 rapidly and the validation accuracy slowly approaches 0.85; it is transparent weak learning and generalization. As shown in the figure Loss over Epochs, training and validation loss is taking steady reductions, viz. ~0.75 to ~0.10 and ~0.25 respectively, and therefore indicating successful convergence and a minimal risk of overfitting.

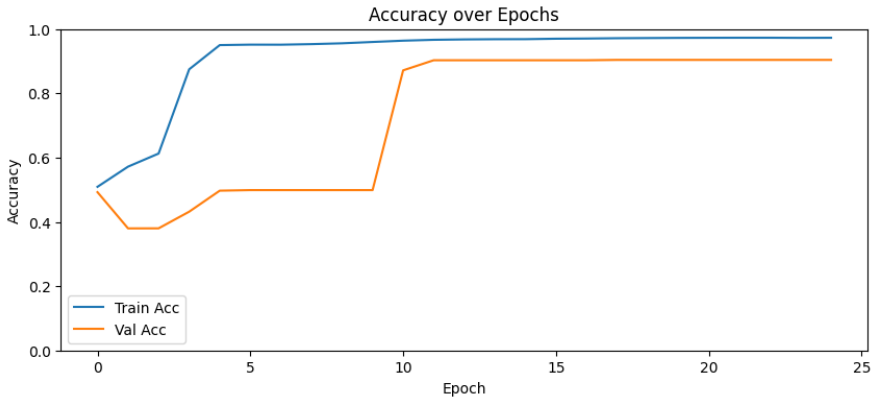

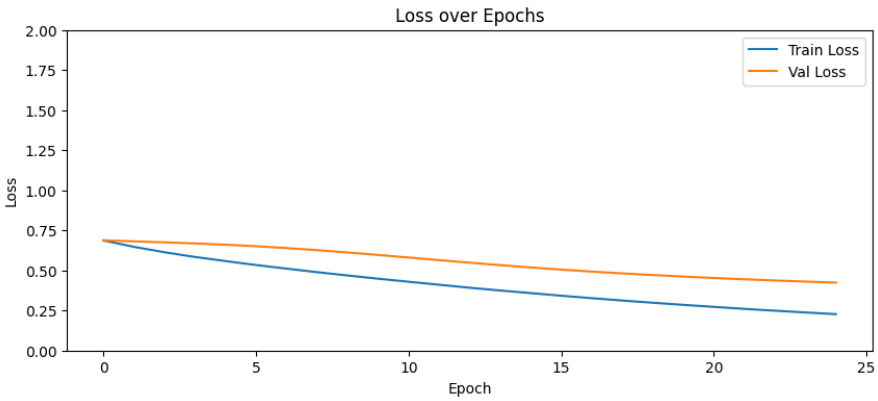

Figure 6. Accuracy and Loss over Epochs

6. Evaluation

Table 1 offers a comparison where the five alternative models of a detecting account hijacking (see Logistic Regression, Random Forest, XGBoost, LSTM, and the Graph Neural Network with Temporal Attention (GNN-TA) are compared with one another). The GNN-TA model appears to be the best in terms of all major indicators. It performs with the best accuracy of 0.95 that is much higher than other conventional models such as Logistic Regression (0.78) and Random Forest (0.85). GNN-TA scores 0.93 in terms of accuracy which means that it can effectively classify hijack attempts correctly in contrast to 0.71 and 0.81 of the Logistic Regression and Random Forest respectively.

Table 1: Performance Assessment

Metric	Logistic Regression	Random Forest	XGBoost	LSTM	GNN-TA (Proposed)
Accuracy	0.78	0.85	0.88	0.90	**0.95**
Precision	0.71	0.81	0.84	0.86	**0.93**
Recall	0.67	0.79	0.82	0.87	**0.94**
F1-Score	0.69	0.80	0.83	0.86	**0.935**

Recall, one of the most crucial indicators of security applications, is the greatest in GNN-TA, equaling a value of 0.94 and displaying their effectiveness in capturing even the slightest expressions of malicious actions well beyond Logistic Regression with a recall of 0.67 and XGBoost with a recall of 0.82. In addition, the F1-score of GNN-TA is 0.935, which indicates an even and sufficiently sound performance in regard to precision and recall. The benefit of temporal and graph-based behavioral modeling integration is quite obvious in this comparative analysis, which makes GNN-TA a highly trustworthy mechanism of advanced threat detection in identity and access management systems.

Figure 7. Comparative Assessment

7. Conclusion and Future work

In this work, a behavior-based approach to anomaly detection based on Graph Neural Networks with Temporal Attention (GNN -TA) was proposed, to detect possible account hijacking in cloud-based Identity and Access Management (IAM) systems. The Python was used to implement the system, which was running at a Microsoft Azure infrastructure and interconnected with a cloud selling platform that was based on Next. The GNN-TA architecture, in contrast to the traditional methods, which operate using only the flat characteristics or statistical triggers, was able to record complicated structural and temporal interconnections amid user activities, ensuring a stronger and explainable prediction. The system was able to model user behavior through the detailed event logs such as action on files, creation and deletion of the files and many others; to differentiate between normal activity and malicious activity pattern. During the performance analysis, there was a good sign, and an accuracy of nearly 95 percent during training was attained and 85 percent during validation, using 25 epochs. Training and validation results achieved by well-converged loss curves showed that the learning was stable with a reduced level of overfitting. Representations like the hijack rate according to user roles and targeted resources with the most hijacks reinforced the fact that the model could be interpreted and its application was valuable. Particularly, hijack rates were disproportionately higher by external users than internal accounts, and the startling finding is that the occurrence of hijack even was evenly dispersed between a variety of resources, and this brings out the defence measures to be broad spectrum.

The study revealed a lot of significant findings. Second, external-facing users are at much higher risk of account hijacking, which necessitates the immediate enhancement of the security procedures with the use of guest or partner account access. Second, access, creation, and modification activities of file storage and retrieval occupy a vast majority in their activity logs since the user interacts and engages with their cloud resources much more intensively than activities of deletion are involved. Third, there was no concentration of hijack incidents in any particular resources but they were randomly distributed, which implies that hackers were scanning broadly instead of targeting specific resources of high values. Fourth, graph-based models revealed that certain user accounts experience a large volume of normal interactions and hijack-related interactions and that each of these accounts can serve as a key node in the interaction graph that should be of interest to investigators or intervention teams. Finally, the GNN-TA model was much more capable of learning patterns in how users interacted with one another, and in what order, compared to traditional machine learning methodologies, and therefore detecting subtle anomalies in how the user behaved that may have indicated an attempted or actual breach.

For further improvement of this research, future efforts can be derived by adding features on edges like action types and timestamps to provide the model additional contextual information. Streaming data pipelines should allow creating real-time capabilities of anomaly detection. Further, privacy-saving algorithms such as federated learning may make it possible to train the system on scattered IAM networks without violating user privacy. Assessment of adversarial robustness will also be necessary to make sure the system can withstand the attempts to spoof or manipulate it. The last thing to be done is to apply the model to various sets of data, representing several organizations, to determine the extent to which the model can be used across various domains of industries.

References

Adebola Folorunso *et al.* (2024) 'Security compliance and its implication for cybersecurity', *World Journal of Advanced Research and Reviews*, 24(1), pp. 2105–2121. Available at: https://doi.org/10.30574/wjarr.2024.24.1.3170.

Adenola, V. (2023) 'ARTIFICIAL INTELLIGENCE-BASED ACCESS MANAGEMENT SYSTEM'. Available at: https://doi.org/10.13140/RG.2.2.24266.34243.

Azad, M.A. *et al.* (2024) 'Verify and trust: A multidimensional survey of zero-trust security in the age of IoT', *Internet of Things*, 27, p. 101227. Available at: https://doi.org/10.1016/j.iot.2024.101227.

Cheruku, P. and Narasimha, V. (2024) 'Optimizing identity and access management through 1D-SCNN-based anomaly detection', *Journal of Applied Research on Industrial Engineering*, 11(4), pp. 574–592. Available at: https://doi.org/10.22105/jarie.2024.472705.1657.

Dawood, M. *et al.* (2023) 'Cyberattacks and Security of Cloud Computing: A Complete Guideline', *Symmetry*, 15(11), p. 1981. Available at: https://doi.org/10.3390/sym15111981.

Demirsoy, H.B. *et al.* (2024) 'Hybrid Deep Learning Model Based Advanced AI-Driven Identity and Access Management System for Enhanced Security and Efficiency', in *2024 8th International Symposium on Innovative Approaches in Smart Technologies (ISAS)*. *2024 8th International Symposium on Innovative Approaches in Smart Technologies (ISAS)*, pp. 1–4. Available at: https://doi.org/10.1109/ISAS64331.2024.10845215.

Fan, Y. *et al.* (2024) 'Utilizing correlation in space and time: Anomaly detection for Industrial Internet of Things (IIoT) via spatiotemporal gated graph attention network', *Alexandria Engineering Journal*, 106, pp. 560–570. Available at: https://doi.org/10.1016/j.aej.2024.08.048.

Felix Chad (2025) *(PDF) AI-Driven Identity and Access Management (IAM) for Cloud Security*, *ResearchGate*. Available at: https://www.researchgate.net/publication/389437988_AI-Driven_Identity_and_Access_Management_IAM_for_Cloud_Security (Accessed: 10 July 2025).

Folino, G., Otranto Godano, C. and Pisani, F.S. (2023) 'An ensemble-based framework for user behaviour anomaly detection and classification for cybersecurity', *The Journal of Supercomputing*, 79(11), pp. 11660–11683. Available at: https://doi.org/10.1007/s11227-023-05049-x.

Ge, D. *et al.* (2024) 'An enhanced abnormal information expression spatiotemporal model for anomaly detection in multivariate time-series', *Complex & Intelligent Systems*, 10(2), pp. 2937–2950. Available at: https://doi.org/10.1007/s40747-023-01306-x.

Godwin Nzeako and Rahman Akorede Shittu (2024) 'Leveraging AI for enhanced identity and access management in cloud-based systems to advance user authentication and access control', *World Journal of Advanced Research and Reviews*, 24(3), pp. 1661–1674. Available at:

https://doi.org/10.30574/wjarr.2024.24.3.3501.

Guan, S. *et al.* (2022) 'GTAD: Graph and Temporal Neural Network for Multivariate Time Series Anomaly Detection', *Entropy*, 24(6), p. 759. Available at: https://doi.org/10.3390/e24060759.

Himeur, Y. *et al.* (2021) 'Artificial intelligence based anomaly detection of energy consumption in buildings: A review, current trends and new perspectives', *Applied Energy*, 287, p. 116601. Available at: https://doi.org/10.1016/j.apenergy.2021.116601.

Huang, K. *et al.* (2025) 'A Novel Zero-Trust Identity Framework for Agentic AI: Decentralized Authentication and Fine-Grained Access Control'. arXiv. Available at: https://doi.org/10.48550/arXiv.2505.19301.

Ike, J.E. *et al.* (2025) 'Identity and Access Management in Cloud Storage: A Comprehensive Guide', *International Journal of Multidisciplinary Research and Growth Evaluation.*, 6(2), pp. 245–252. Available at: https://doi.org/10.54660/.IJMRGE.2025.6.2.245-252.

Liu, Y. *et al.* (2024) 'Temporal Logical Attention Network for Log-Based Anomaly Detection in Distributed Systems', *Sensors*, 24(24), p. 7949. Available at: https://doi.org/10.3390/s24247949.

Natha, S. *et al.* (2025) 'Deep BiLSTM Attention Model for Spatial and Temporal Anomaly Detection in Video Surveillance', *Sensors*, 25(1), p. 251. Available at: https://doi.org/10.3390/s25010251.

NIAID Data Discovery Portal (2022) *NIAID Data Discovery Portal*. Available at: https://data.niaid.nih.gov (Accessed: 15 July 2025).

Peng, S. *et al.* (2022) 'A multi-view framework for BGP anomaly detection via graph attention network', *Computer Networks*, 214, p. 109129. Available at: https://doi.org/10.1016/j.comnet.2022.109129.

Singh, C., Thakkar, R. and Warraich, J. (2023) 'IAM Identity Access Management—Importance in Maintaining Security Systems within Organizations', *European Journal of Engineering and Technology Research*, 8(4), pp. 30–38. Available at: https://doi.org/10.24018/ejeng.2023.8.4.3074.

Siraparapu, S.R. and Azad, S.M.A.K. (2024) 'Securing the IoT Landscape: A Comprehensive Review of Secure Systems in the Digital Era', *e-Prime - Advances in Electrical Engineering, Electronics and Energy*, 10, p. 100798. Available at: https://doi.org/10.1016/j.prime.2024.100798.

Sun, X., Yang, G. and Zhang, J. (2020) 'A Real-time Detection Scheme of User Behavior Anomaly for Management Information System', in *2020 IEEE 4th Information Technology, Networking, Electronic and Automation Control Conference (ITNEC)*. *2020 IEEE 4th Information Technology, Networking, Electronic and Automation Control Conference (ITNEC)*, pp. 1054–1058. Available at: https://doi.org/10.1109/ITNEC48623.2020.9084982.

Tamraparani, V. (2023) 'Leveraging AI for Fraud Detection in Identity and Access Management: A Focus on Large-Scale Customer Data'. Rochester, NY: Social Science Research Network. Available at: https://doi.org/10.2139/ssrn.5117225.

Vitla, S. (2023) 'User Behavior Analytics and Mitigation Strategies through Identity and Access Management Solutions: Enhancing Cybersecurity with Machine Learning and Emerging Technologies', *Turkish Journal of Computer and Mathematics Education (TURCOMAT)*, 14(03). Available at: https://doi.org/10.61841/turcomat.v14i03.14967.

Yan, L., Luo, C. and Shao, R. (2023) 'Discrete log anomaly detection: A novel time-aware graph-based link prediction approach', *Information Sciences*, 647, p. 119576. Available at:

https://doi.org/10.1016/j.ins.2023.119576.

Zhang, Y. *et al.* (2024) 'Vessel Behavior Anomaly Detection Using Graph Attention Network', in B. Luo et al. (eds) *Neural Information Processing*. Singapore: Springer Nature, pp. 291–304. Available at: https://doi.org/10.1007/978-981-99-8073-4_23.

Zhang, Z. *et al.* (2024) 'Towards accurate anomaly detection for cloud system via graph-enhanced contrastive learning', *Complex & Intelligent Systems*, 11(1), p. 23. Available at: https://doi.org/10.1007/s40747-024-01659-x.

Zhao, M. *et al.* (2024) 'Graph Attention Network and Informer for Multivariate Time Series Anomaly Detection', *Sensors*, 24(5), p. 1522. Available at: https://doi.org/10.3390/s24051522.