AI-DRIVEN PRIVILEGED ACCESS MANAGEMENT (PAM): A CONTEXT-AWARE DEEP LEARNING FRAMEWORK FOR PREDICTIVE ANOMALY DETECTION IN ENTERPRISE SYSTEMS

Anupam Nandan
Senior Manager Cybersecurity Consulting
Ernst and Young
realanupamnandan@gmail.com

To Cite this Article
Anupam Nandan, "AI-Driven Privileged Access Management (PAM): A Context-Aware Deep Learning Framework for Predictive Anomaly Detection in Enterprise Systems", Journal of Science Engineering Technology and Management Science, Vol. 03, Issue 07, July 2026, pp: 781-790, DOI: http://doi.org/10.64771/jsetms.2026.v03.i07.pp781-790
Submitted: 07-06-2026 *Accepted: 14-07-2026* *Published: 20-07-2026*

***Abstract*—Privileged Access Management (PAM) systems play a critical role in securing high-risk identities within enterprise environments. Current PAM solutions, however, are mostly based on static policies and context-independent monitoring, which are not enough to detect context-dependent anomalous behaviors in contemporary distributed systems. In this paper, a context-aware, AI-based anomaly detection mechanism based on hybrid deep learning architectures to learn temporal and behavioral patterns of privileged access activities is proposed. The proposed CNN-BiLSTM-Attention network, Transformer-GRU fusion network, Attention-CNN-LSTM network and Hierarchical BiGRU with context gating network are integrated together. Experimental results on a synthetic enterprise PAM dataset show the remarkable detection performance of the Hierarchical BiGRU with an accuracy of 99.14% and an F1 score of 99.32%. The work extends PAM from a static-to-dynamic approach to access control to a more continuous approach to security monitoring.**

Keywords*—Privileged Access Management, Identity and Access Management, anomaly detection, deep learning, contextual risk scoring, explainable AI, enterprise cybersecurity.

I. INTRODUCTION

The attack surface for enterprise cybersecurity threats has grown larger due to the constant changes in enterprise IT environments that have come with the move to cloud adoption, distributed architectures, and greater importance placed on machine identities [1]. One of the most prominent risk vectors is privileged accounts that give access to sensitive systems, infrastructure and data that are elevated [2]. By implementing access control, credential management, and session monitoring, PAM solutions can help reduce these risks. Indeed, practical implementations of PAM in large enterprises such as financial services and healthcare illustrate a major drawback of PAM: It can effectively control access before authentication, but has less success in detecting misuse after access is granted [3][4]. In larger-scale environments, privileged access may be long-lived, context-independent and not necessarily related to the user's intent[5].

In regulated industries, these are common issues that have been seen in practice: Over-accumulation of privileges, lack of visibility of behavioural anomalies, finding insider threats is difficult, and rule-based monitoring yields a lot of false positives. From the operational experience, the following issues have been observed: Excessive privileges being accumulated, limited visibility of behavioural anomalies, difficulty in identifying insider threats and misuse of credentials, and high false positive rates in rule-based monitoring [6]. However, traditional methods have a limited ability to detect temporal patterns, behavioural anomalies, and context such as device, location, role and resource sensitivity [7].

To overcome this limitation, context-aware anomaly detection: first, looks beyond individual events to the context, i.e., the surrounding behaviour and environment. While analyzing access events is important, contextual analysis also takes into account the role of the user, the time of the access, the type of device, the geographic location, the type of action and the sensitivity of the resource to determine if there is any deviation from the expected patterns of privileged access [8].

AI and ML can be used to build access logs at scale and identify slight variations in privileged access logs [9]. But these approaches do not necessarily feature contextual awareness or sequential modelling, and security operations are often not interpretable. Some of the existing studies also deal with the generic network traffic analysis or generic system log anomaly detection, instead of behavioural analytics techniques in enterprise PAM environment [10][11]. This paper aims to fill the above gaps by proposing an Anomaly Detection framework for PAM system based on context-awareness.

The proposed framework will be an AI-informed behavioural analytics layer above the telemetry provided by the PAM, and is not meant to be a substitute for commercial PAMs. The framework adopts hybrid sequential deep learning architectures such as CNN-BiLSTM-Attention, self-attention-based GRU fusion, Attention-CNN-LSTM and Hierarchical BiGRU with context gating to capture the temporal and behavioural patterns in privileged access activities. The framework also incorporates SHAP-

based explainability, contextual risk scoring, anomaly visualization and alert prioritization, which are aimed to facilitate enterprise security monitoring workflows:

- Proposes a context-aware framework of PAM anomalies detection for the analysis of privileged access risk combining temporal, behavioural and environmental signals.
- Develops and tests hybrid deep learning models such as CNN-BiLSTM-Attention, Transformer-GRU fusion, Attention-CNN-LSTM and Hierarchical BiGRU with context gating.
- Implements enterprise-informed feature engineering using large-scale deployments in financial and healthcare industries, whilst keeping any client-specific data confidential.
- Embeds SHAP into the model to find the features that impact model decisions, and builds operational trust.
- Integrates powerful deep learning results with complementary anomaly detection, risk scoring, alert prioritization and visualization.

A. Justification and Novelty

The conventional Privileged Access Management (PAM) systems primarily focus on authentication, monitoring based on rules, and static access policies, which can often be ineffective against detecting contextually privileged misuse, insider threats or atypical behaviour after logon. Intelligent behavioural monitoring is needed to analyze temporal access patterns, contextual deviations and sequential privileged activities, which are required in modern enterprise environments. In this paper, a behavioural anomaly detection framework utilizing AI to provide context-awareness is proposed to overcome these limitations. The novelty of the proposed work is the combination of hybrid sequential deep learning architectures, contextual behavioural analytics, explainable AI, multi-factor risk scoring with intelligent alert prioritization specifically for the PAM telemetry monitoring. The proposed framework dynamically learns privileged user behaviour, detects contextual anomalies, enhances insider threat visibility, and enables adaptive enterprise monitoring of privileged users with explainable, risk-aware behaviour analytics, unlike traditional PAM monitoring methods.

II. LITERATURE REVIEW

AI, ML and DL have been making great strides in recent times to enhance anomaly detection, behavioural analysis and enterprise cybersecurity monitoring systems.

E. Yilmaz et al. (2026) address the conceptual concepts and life cycle needs of machine identities and summarizes the latest research for certificate and key management in distributed environments. A comprehensive analysis of strengths and weaknesses of the existing PAM implementations and potential research directions is identified. Thus, Machine identity governs is crucial for the security and operational resiliency of today's enterprise systems [12].

S. K. Udayakumar et al. (2025) proposed an ensemble anomaly detection framework of QIC, Autoencoder, Graph Learning, HMM, Hyperdimensional Computing and Recurrence Quantification (RQ) to improve the anomaly detection in enterprise environment. The ensemble model outperforms the traditional ML and DL algorithms, with an almost 90% accuracy, 94% precision, 92% recall and 93% F1-score [13]. Likewise, C. Vijitphu and N. Utakrit (2025) provides anomaly and data breach detection system by analyzing logs of internet traffic using Python. This is a Power BI dashboard, which is presented on a web window to enable executives and security administrators to get actionable insights. The accuracy was validated by the experts and was 90.67% and user satisfaction was obtained with high rate [14]

H. Uchida et al. (2024) present a novel method for detecting parameters' anomalies using unsupervised learning using BertForMaskedLM. BertForMaskedLM is a model that makes predictions for masked parts in the input string and outputs the probabilities. The authors obtained a good performance in terms of F1-score for several anomaly patterns [15] by this method. Also, ML techniques have been applied to benefit financial transaction monitoring. Financial transaction monitoring has also benefited from ML techniques. B. Torky et al. (2024) make use of the ML and AI capabilities to augment the payment transaction anomaly detection with the rule-based approach using SARIMAX, Facebook Prophet and SVM models. The SVM model was the one that was evaluated in the experiment and it has shown to give the best predictive performance [16]

There are further improvements in anomaly detection of system logs, thanks to the development of representation learning. S. Yan et al. (2023) have presented a contrastive self-supervised representation learning-based method for log anomaly detection. This method eliminates the error that can occur with regular log parsing and makes the detection process easier. Experimental results on the HDFS dataset showed an F1-score of 96.7%, better than a few baseline methods, such as supervised and unsupervised methods [17].

Research Gap: Previous work has discussed anomaly detection using machine learning and deep learning in various network security, system logs, cloud-based and financial fraud monitoring scenarios. The methods that have proven to be successful in identifying abnormal activity patterns include ensemble methods, self-supervised learning, transformer-based models and graph-based techniques. Most of the existing literature, however is on generic log anomaly detection or network-level monitoring, and not on specific use cases of the PAM. PAM introduces unique challenges where privileged identities have high risk exposure, access events are high-risk even if they are rare, and behavioural context is as much a part of the access activity as is the access activity itself. This work builds on previous research to apply a hybrid sequential model, contextual feature engineering, explainable AI and contextual risk scoring specifically for privileged access behaviour and PAM telemetry monitoring in enterprise applications. The overall proposed AI-assisted PAM telemetry system and anomaly detection flow are shown in Fig. 1

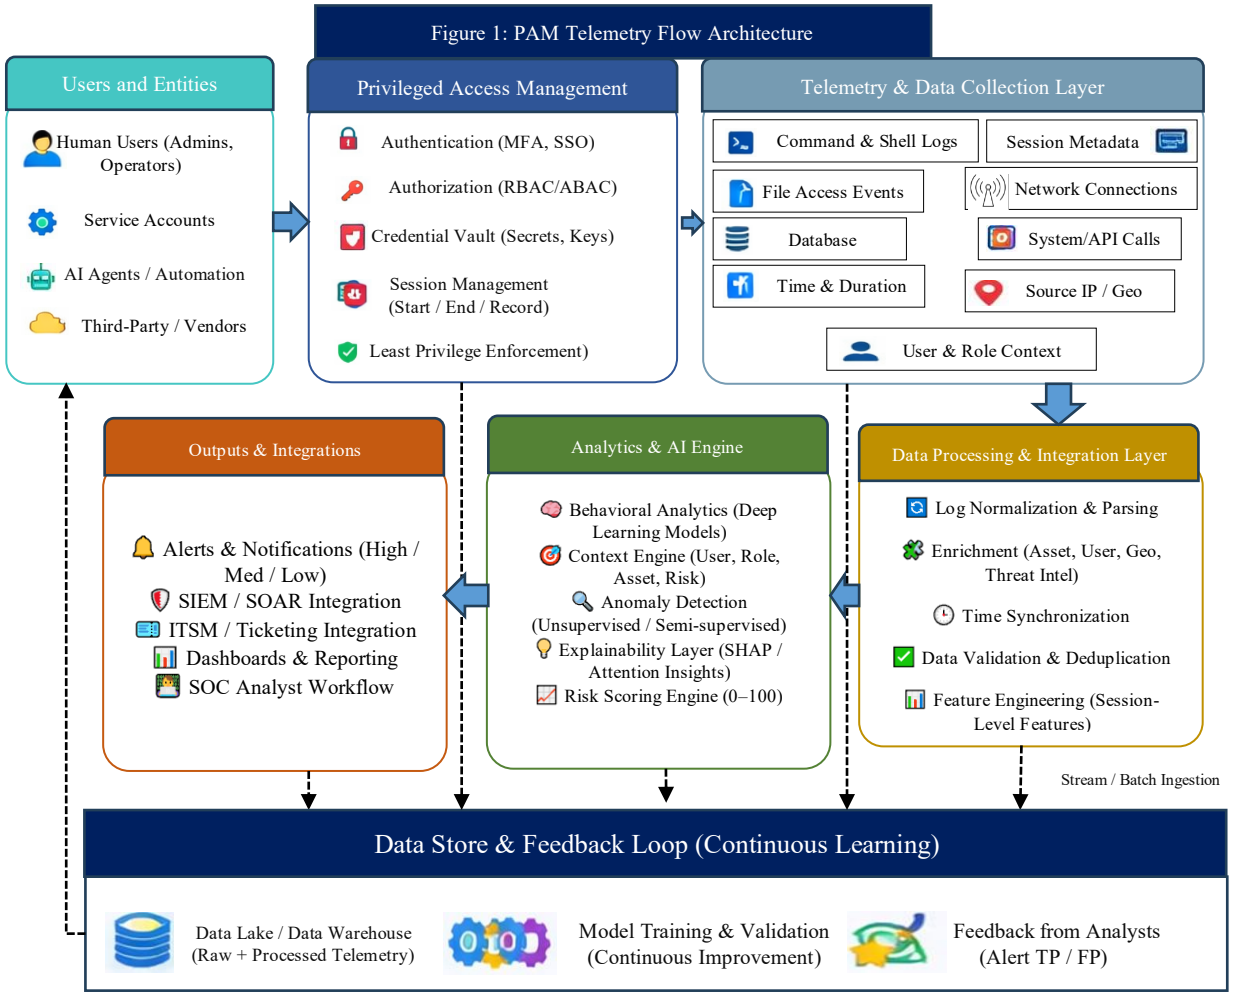

Fig. 1. Proposed AI-Driven PAM Telemetry Architecture for Context-Aware Behavioural Anomaly Detection

III. METHODOLOGY

The proposed model introduces an approach for anomaly detection based on behavioural analytics in enterprise privileged access monitoring systems that considers the entity's context. Synthetic PAM telemetry logs are used for behavioural analysis and include user behaviours, access events, device information, location information and temporal access information. First, preprocessing of the data and feature engineering based on contextual PAM are done, and the temporal patterns and behaviour of accesses are extracted. Contextual behavioural sequences are then prepared through additional attribute encoding using context-aware, data balancing with SMOTEENN and feature scaling, as well as through modelling by accessing to the behaviour sequences sequentially with PAM.

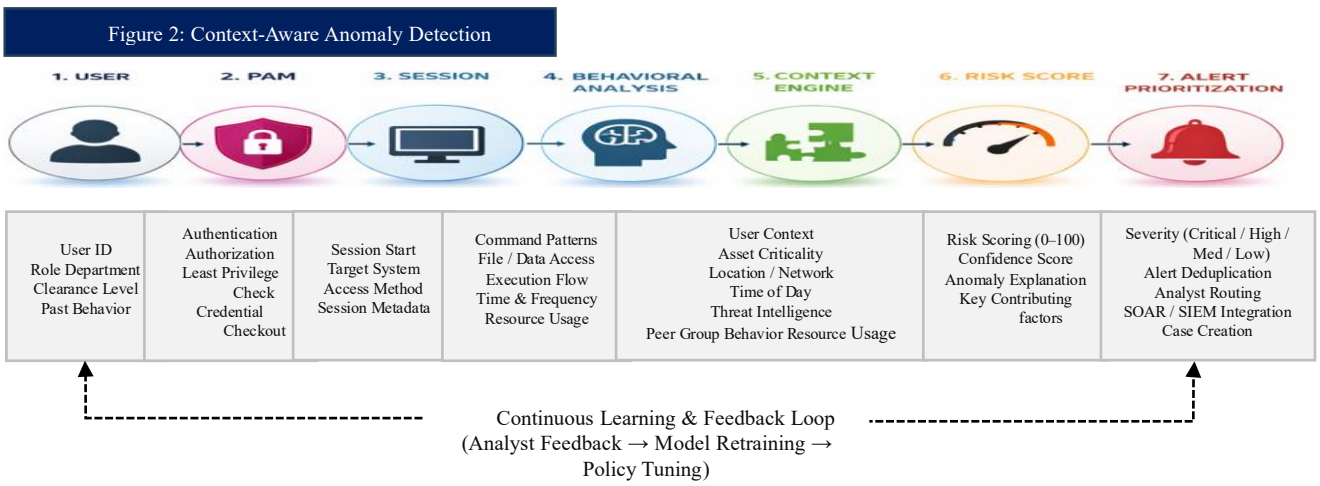

Fig. 2. Context-Aware Deep Learning Workflow for Privileged Access Anomaly Detection

The framework employs hybrid sequential deep learning architectures that leverage CNN-BiLSTM-Attention, Transformer-GRU-Dense Attention Fusion, Attention-Augmented CNN-LSTM, and Hierarchical BiGRU-Attention models to effectively support and learn temporal privileged access dependency and contextual behavioural deviations. Additional enhancements include abnormal access behaviour, device deviations, location changes and off-hours activities analyzed by privileged user behavioural profiling and context-aware PAM risk scoring. Unsupervised anomaly validation is possible with the support of auxiliary behavioural anomaly detection models, based on Autoencoder and Isolation Forest models. Moreover, Explainable AI methods such as SHAP and LIME enhance model interpretable, and visualization using UMAP aids in understanding behavioural anomaly clustering. Lastly, there's the intelligent PAM alert generation and prioritization framework – which classifies suspicious privileged access activities into various severity levels to enable real-time monitoring of enterprise security events and behavioural threat analysis. The overall context-aware anomaly detection and behavioural analytics process incorporated in the proposed system is shown in Fig. 2.

A. Data Gathering

This study is based on a simulated Privileged Access Management (PAM) enterprise-level telemetry dataset which should represent access scenarios in the enterprise. It includes around 50,000 records and 16 contextual elements such as user identities, the time they accessed the resource, their location, device type, user roles, and resource usage patterns that are commonly seen in enterprise PAM audit logs. The dataset mimics access patterns that are realistic, such as access by roles, logins from distant locations, and access patterns over time. Real enterprise PAM logs are not easily available for public use, because of security and/or privacy concerns, so synthetic PAM telemetry datasets are becoming a common source for enterprise security research and anomaly detection modelling. The behaviour trend, context relations and feature correlations are further analysed using exploratory data analysis and visualization techniques for intelligent telemetry monitoring and control of PAMs.

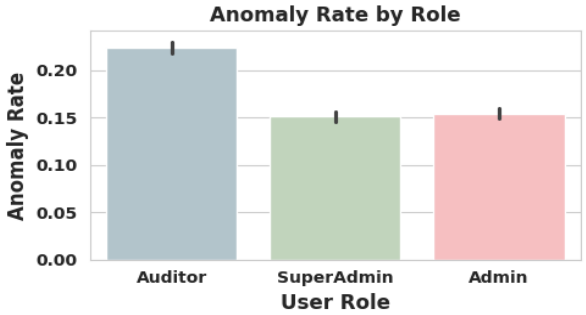

Fig. 3. Anomaly Rate Comparison Across User Roles

In Fig 3, the anomaly distribution of the privileged user roles is shown. The highest anomaly rate is for the Auditor role at 22%, with the SuperAdmin and Admin roles having almost 15–16% anomaly rates. The results indicate the need for role-aware behavioural monitoring on enterprise PAM environments..

Anomaly probabilities are shown in Fig. 4 for the various combinations of location and device types and the various types of resources. The higher anomaly rates are seen in case of Network_Device (21.4%) and Server (21.1%) resources as compared to Application and Database resources. Furthermore, remote mobile access has the highest percentage of anomalies at 38.82%, indicating that there is a higher behavioural risk in enterprise PAM environments[18].

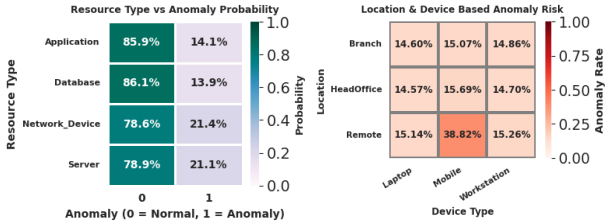

Fig. 4. Heatmap Analysis of Anomaly Probabilities and Risks by Resource Type and Location/Device

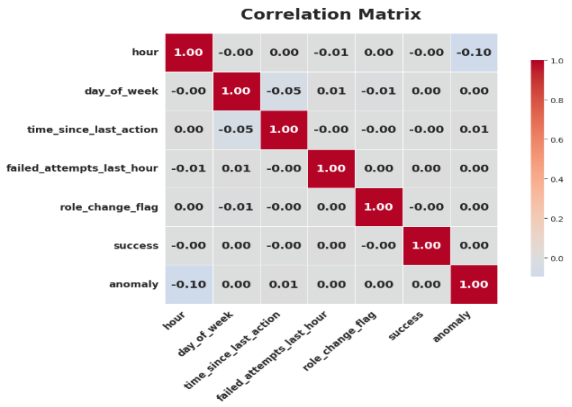

Fig. 5. Correlation Matrix of Numerical Features

The correlation matrix of the contextual and behavioural features in the PAM telemetry dataset is presented in Fig. 5. The correlation values of most of the features are low, showing low degree of redundancy among features. Furthermore, the anomaly attribute correlates with individual attributes in a weak manner, indicating that an anomalous privileged behaviour does not depend on individual features, but rather a bunch of sequential and contextual attributes.

B. Data pre-processing

Unprepared data needs to be pre-processed before effective detection of behavioural anomalies and sequential modelling of the PAM telemetry data can be done. Access timestamps are used to derive temporal attributes, such as hour, weekday, month, weekend activity and off-hours behaviour to capture the context of access. Cyclic sin and cos encoding is also used on temporal features, to maintain sequential temporal relationships. Label encoding converts categorical attributes such as user role, device type, location or interactions with the resource, into numerical values. After the processing operation, the processed data set is split into 80% training and 20% testing samples, keeping the same anomaly distribution as the original data set so that the behaviour anomaly evaluation can be done in an unbiased manner. The final data set includes features of enterprise privileged access environments and activities, both under normal and abnormal contexts and temporality.

C. Data Balancing

The dataset for the PAM shows a class imbalance, indicating that privileged access is an abnormal event, and is not as common as normal behaviour. The SMOTEENN balancing technique is used to enhance the performance of anomaly detection and learning of behavioural patterns with these techniques. The anomaly detection performance and behavioural pattern learning are enhanced by applying these techniques to the training data set using the SMOTEENN balancing technique. SMOTEENN is a hybrid minority oversampling and noise reduction method to improve the class representation and remove overlapping behavioural patterns.

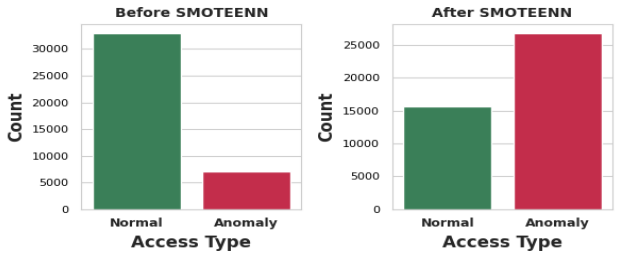

Fig. 6. Anomaly Distribution Before and After SMOTEENN

The distributions of the anomalies are shown before and after SMOTEENN in Fig. 6. The balancing process enhances the learning of the abnormally privileged activities which allows the proposed framework to learn both normal and abnormal behaviours in enterprise PAM environments.

D. Scaling

To ensure the normalisation of the features/attributes in the dataset, data normalisation is done to ensure uniformity in the scale of the numerical features, improving the stability of the model in the process of behavioural anomaly detection. Less sensitive to outliers seen in enterprise PAM telemetry data, the RobustScaler technique is used. The feature representation is "normalized" using Equation. (1):

$$x_i' = \frac{x^i - Q_2(x)}{Q_3(x) - Q_1(x)} \qquad (1)$$

where Q_1(x), $Q_2(x)$, and $Q_3(x)$ are the first quartile, median and third quartile of the feature (x), respectively..

E. Proposed Model

The proposed PAM behavioural analytics framework uses multiple hybrid sequential deep learning architectures to analyze privileged access behaviour and the patterns of anomalous situations.

1) Transformer-GRU-Dense Attention Fusion

The Transformer-GRU-Dense Attention Fusion model takes as input sequences of shape (5, 21), where 5 is the number of images and 21 is the number of features, and uses 8 heads of self-attention to extract contextual and privileged access behaviour of the same length from the input sequence using stacked GRU layers of size 128 and 96. To further boost anomaly classification, the use of context attention and global pooling as well as dense layers of 256 and 128 neurons, is employed. The model has around 312K trainable parameters, and it uses the AdamOptimizer (LearningRate = 0.001) to optimize the model using the following evaluation metric: binary cross-entropy loss, accuracy, and AUC.

2) Hybrid CNN-BiLSTM-Attention

The hybrid CNN-BiLSTM-Attention model consists of two 1D convolution layers (128 and 96) and two Bidirectional LSTM layers (128 and 64) and an 8-head self-attention mechanism for sequential behavioural anomaly detection. Anomaly classification is done using global pooling and dense layers of 256 and 128 neurons respectively. The model has around 806K of parameters and was trained via Adam optimization, binary cross-entropy loss, and measured by the AUC.

3) Attention-Augmented CNN-LSTM with Residual

It features multi-scale convolutional layers, parallel LSTM layers (128, 96 units), self-attention, residual connection and dense attention for contextual feature learning. The robustness of behavioural anomaly detection is enhanced for PAM telemetry environments with global pooling and dense layers of 256, 128 and 64 neurons with normalization and dropout.

4) Hierarchical BiGRU-Attention with Context Gating

The Hierarchical BiGRU-Attention model consists of stacked Bidirectional GRU models with 128, 96, and 64 units and an 8-head Multi-Head Attention mechanism, which is able to model hierarchical temporal dependencies in privileged access behaviour. Context gating, residual connections, global pooling and intermediate dense layers (256, 128, and 64) are further improvements of the contextual anomaly representation. The model has approx. 1.02M trainable parameters and is trained with AdamOptimizer (LearningRate = 0.001) loss function as binary cross-entropy and data validation method is stratified cross-validation.

F. Evaluation Parameters

The performance of the proposed behavioural anomaly detection models is evaluated using a confusion matrix in enterprise PAM environment. The matrix provides a summary of expected and actual class label values based on the results of True Positive (TP), False Positive (FP), True Negative (TN), and False Negative (FN). The evaluation metrics such as Accuracy, Precision, Recall, F1-Score, AUC-ROC and MCC are calculated based on these values for assessing the performance of anomaly detection. Equations (2) to (7) are the evaluation metrics corresponding to the above equations, respectively.

$$Accuracy = \frac{\mathrm{TP+TN}}{\mathrm{TP+Fp+TN+FN}} \quad (2)$$

$$Precision = \frac{\mathrm{TP}}{\mathrm{TP+FP}} \quad (3)$$

$$Recall = \frac{\mathrm{TP}}{TP+FN} \quad (4)$$

$$F1-score = 2 \times \frac{Precision \times Recall}{Precision+Recall} \quad (5)$$

$$Specificity = \frac{TN}{TN+FP} \quad (6)$$

$$MCC = \frac{TP \cdot TN - FP \cdot FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}} \quad (7)$$

The accuracy shows how accurate the anomaly predictions are, precision and recall show how well the model can correctly predict anomalous privileged access activities. F1 Score is a balanced measure that combines precision and recall. Moreover, in imbalanced anomaly detection scenarios, MCC is a reliable evaluation; and AUC-ROC is used to measure the model's ability to differentiate between normal and anomalous privileged behaviour.

G. Model Interpretability

In enterprise PAM, the SHAP and LIME methods are used to pinpoint key contextual and behavioural factors that impact deep learning models' predictions of anomalies, enhancing interpretability. Also, UMAP visualization is applied for anomaly detection on separation of normal and anomalous privileged access patterns.

IV. Results and Discussion

The experiments are performed on the Windows 10 system equipped with Intel Core i7-10750H CPU (2.60 GHz, 6 cores) and 16 GB RAM with the help of the Jupyter Notebook environment in the Anaconda platform. This section shows the performance of the proposed hybrid behavioural anomaly detection models and compares it with auxiliary anomaly detection models. Table I shows the performance of all the proposed models with the Hierarchical-BiGRU model showing the best overall performance.

TABLE I. PERFORMANCE METRICS COMPARISON OF DEEP LEARNING MODELS FOR ANOMALY DETECTION

Model	Accuracy	Precision	Recall	F1-Score	AUC-ROC	MCC	Balanced Accuracy	Specificity
CNN-BiLSTM-Attention	98.85	98.62	99.57	99.09	0.998691	97.55	98.60	97.63
Transformer-GRU-Fusion	98.98	98.17	99.77	98.96	0.997851	97.19	98.30	96.82
Attention-CNN-LSTM	98.50	98.66	98.96	98.81	0.998117	96.78	98.33	97.71
Hierarchical-BiGRU	99.14	98.94	99.71	99.32	0.999151	98.16	98.84	98.17

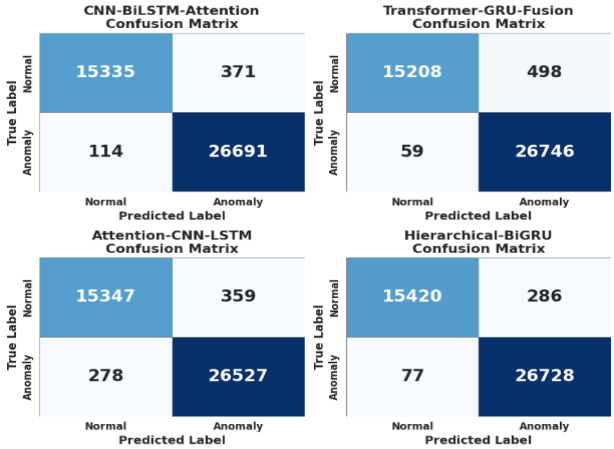

Fig. 7. Confusion Matrix Comparison of Proposed Hybrid Behavioural Models

The privileged access behavioural anomaly detection performance of all proposed hybrid models is compared to each other in Fig. 7. The Hierarchical-BiGRU model outperformed the other models in terms of the classification accuracy, the lowest misclassification rates and best ability to identify anomalies. In general, the results show how well the hybrid sequential architectures can capture the contextual privileged access behaviour and temporal anomaly patterns in an enterprise PAM telemetry environment.

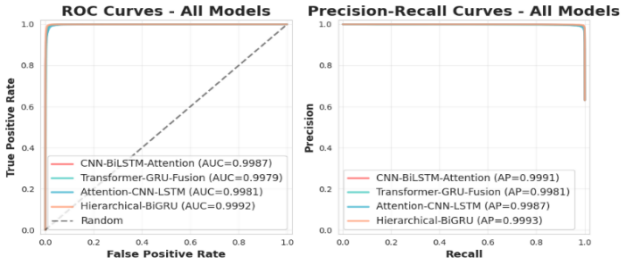

Fig. 8. ROC and Precision-Recall Curve Comparison of Proposed Hybrid Models

The comparison of the ROC and Precision-Recall curve of the proposed hybrid behavioural models is shown in Fig. 8. They were nearly perfect in their classification performance (with high average precision and AUC) in the ability to differentiate normal and anomalous privileged access behaviour in all architectures. The best overall separability and predictive reliability in the PAM contextual behavioural anomaly detection was obtained from the Hierarchical-BiGRU model.The Hierarchical-BiGRU model outperformed the other models in terms of the overall separability and predictive reliability for contextual PAM behavioural anomaly detection.

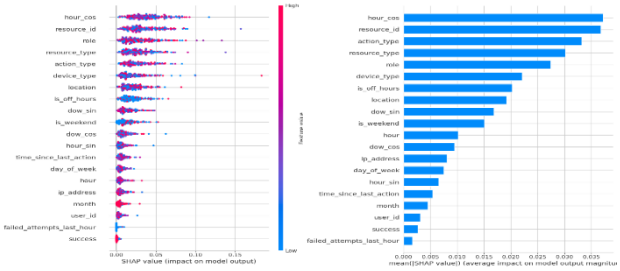

Fig. 9. SHAP-Based Feature Importance Analysis for Contextual PAM Behavioral Modeling

Fig 9 shows the feature importance analysis conducted by using the SHAP values of the proposed behavioural anomaly detection framework. This analysis reveals that hours of the day (hour_cos), resource (resource_id), action type (action_type), resource type (resource_type), and role (role) are all key temporal, resource and contextual behavioural features that play a significant role in the prediction of anomalies. The analysis helps to elucidate, in a better way, the most influential telemetry features of the context associated with abnormal privileged access to models.

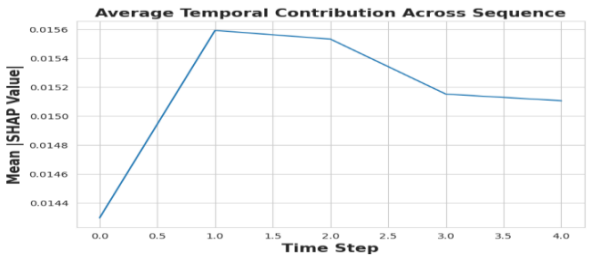

Fig. 10. Temporal Contribution Analysis Across Sequential PAM Access Patterns

The mean SHAP values are used for the illustration of the temporal contribution of sequential access events that are shown in Fig 10. The findings suggest that the earlier time steps are more important in the context of anomaly prediction, therefore, the temporal behavioural dependencies are crucial in the context of contextual PAM telemetry analysis as well as to privileged access anomaly detection.

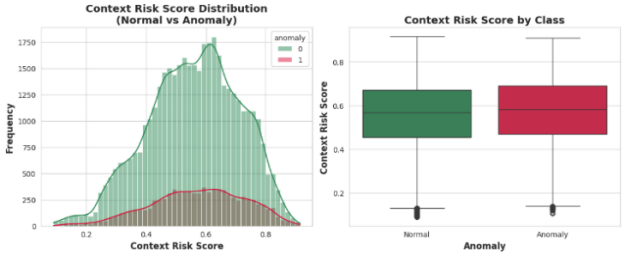

Fig. 11. Class-wise Context Risk Score Distribution

Fig 11 shows a distribution of the context-aware risk scores for normal and anomalous privileged access activities. As can be seen from the results, it is generally observed that the anomalous behaviours have higher contextual risk scores which highlights the effectiveness of the proposed context-aware privileged access anomaly risk (PAM) risk scoring mechanism in differentiating privileged access anomalies..

A. Auxiliary Behavioural Anomaly Validation

The contextual behavioural anomaly validation is then extended with the help of auxiliary unsupervised anomaly detection with Autoencoder and Isolation Forest models in enterprise PAM environments. Both models showed low performance as compared to the proposed hybrid sequential architectures, but they were effective in identifying abnormal Deviations from privileged access: Reconstruction error & anomaly score analysis.

TABLE II. RESULTS OF AUTOENCODER AND ISOLATION FOREST FOR BEHAVIOURAL ANOMALY DETECTION

Metric	Autoencoder	Isolation Forest
Accuracy	0.7736	0.7081
Precision	0.1128	0.1531
Recall	0.0414	0.1446
F1-Score	0.0606	0.1487
AUC-ROC	0.4307	0.4579
MCC	-0.0438	-0.0273

In order to compare the Auxiliary Behavioral anomaly detection performance of the Autoencoder and the Isolation Forest models, their performance is compared in Table II. The performance of anomaly identification by using Isolation Forest turned to be comparatively good, and the capability of reconstruction-based behavioural learning by using Autoencoder was comparatively steady.

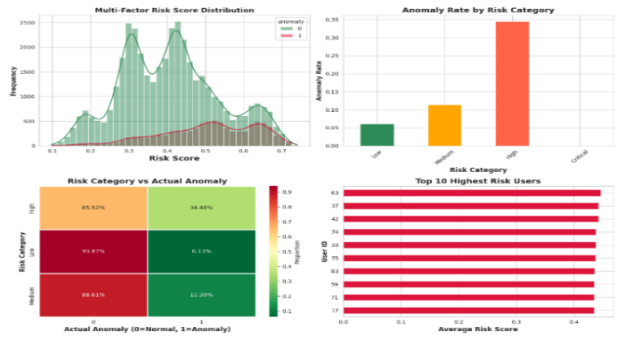

Fig. 12. Multi-Factor PAM Risk Score and Behavioural Risk Analysis

Fig 12 shows the PAM multi-factor risk scoring analysis of contextual privileged access monitoring. As can be seen in the results, the anomalous privileged activities are more focused on high-risk group, and the normal activities are mostly included in

low and medium risk groups. The analysis also emphasizes the ability of the proposed framework to detect high-risk privileged users and enables intelligent threat prioritization for enterprise PAM use.

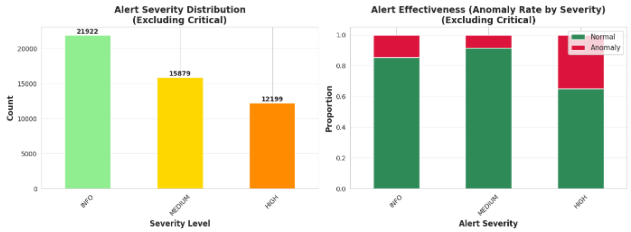

Fig. 13. Alert Severity Distribution and Effectiveness Analysis

The distribution of the proposed framework of prioritization of alerts for PAM is shown in Fig 13, and the effectiveness of the proposed framework is shown as presented in the figure. The outcomes reveal that the severity of the alerts correlates directly with the likelihood of privileged access anomaly; the higher the severity level, the greater the association with the likelihood of privileged access anomaly.

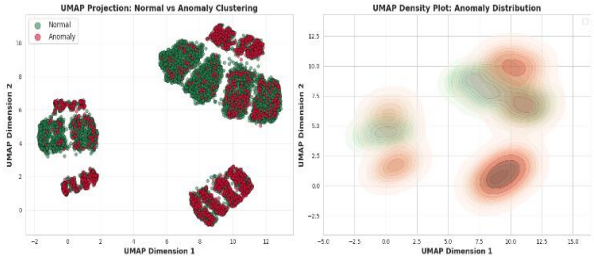

Fig. 14. UMAP-Based Behavioural Anomaly Clustering Visualization

The distribution of the proposed framework of prioritization of alerts for PAM is shown in Fig 14, and the effectiveness of the proposed framework is shown as presented in the figure. The outcomes reveal that the severity of the alerts correlates directly with the likelihood of privileged access anomaly; the higher the severity level, the greater the association with the likelihood of privileged access anomaly.

B. Enterprise Implication

The framework proposed is based on experience gained from applying the PAM in the regulated financial services and healthcare sectors. These environments create many challenges for privileged access governance due to the number of users, distributed infrastructure, complex application portfolios and high audit expectations. A behaviour-based model can help increase visibility into the misuse that happens once legitimate authentication. The framework enables real-world security results like fewer false positives using context-aware scoring, prioritization of high-risk privileged activity, better detection of insider misuse, and more monitoring control evidence for ongoing monitoring. The framework combines multiple architectures of hybrid deep learning, contextual behavioural profiling, explainable AI, multi-factor risk scoring, and intelligent alert prioritization to improve real-time privileged access monitoring and contextual anomaly analysis in enterprise privileged access monitoring (PAM) platforms.

TABLE III. COMPARISON WITH EXISTING PAM SOLUTIONS

Feature	**CyberArk**	**BeyondTrust**	**Proposed AI-PAM System**
Risk-Based Authentication	Rule-Based	Policy-Based	AI-Based
Behavioral Analysis	Limited	Moderate	Advanced (User Profiling)
AI/ML Integration	Partial	Partial	Full (Deep Learning + ML)
Real-Time Threat Detection	Yes	Yes	Yes (Real-Time Engine)
Just-In-Time Access (JIT)	Yes	Yes	Yes (Dynamic JIT)
Privilege Escalation Detection	Yes	Yes	Yes (AI-Based)
Anomaly Detection	Basic	Basic	Advanced (DL + AE + IF)

Explainable AI (XAI)	No	No	Yes (SHAP)
SIEM Integration	Yes	Yes	Yes
Adaptive Access Control	Limited	Moderate	Fully Adaptive

Table III shows how the proposed AI-based PAM framework is similar and different from other PAM solutions (CyberArk and BeyondTrust). The comparison shows the advanced behavioural analytics, advanced Contextual anomaly detection, advanced Explainable AI integration, Adaptive risk scoring and Intelligent threat prioritization capabilities that go beyond traditional rule-based PAM monitoring approaches. It also serves as a basis for future integration with just-in-time access, zero standing privilege, machine identity management and AI-agent access control..

C. Limitations & Future Work

The major drawback of the study is that a synthetic dataset has been used. As the dataset is meant to mimic enterprise PAM activity, it is essential to validate the data in a production environment to determine production performance, operational limits and environmental variations. Real-time monitoring systems could also suffer from computational costs due to the use of deep learning models. The framework will be considered in future work based on telemetry of enterprise PAM, streaming pipelines and adaptive learning models. Other research needs to continue in the machine identities, service accounts, workload identities, and AI agents, in which the misuse of privileges could happen at machine speed and requires constant authorization decisions. Further improvements can include graph-based behavioural modelling, federated security analytics and continually adaptive anomaly detection systems to enterprise privilege governance.

REFERENCES

[1] K. M. R. Seetharaman and P. Yadav, "A Machine Learning Framework for Detecting and Mitigation of Cyber Threats in IoT Environments," in *2025 3rd International Conference on Inventive Computing and Informatics (ICICI)*, IEEE, Jun. 2025, pp. 1112–1119. doi: 10.1109/ICICI65870.2025.11069697.

[2] M. Bhargude, "Privileged Access Management: Ensuring Security and Accountability," *Int. J. Adv. Res. Sci. Commun. Technol.*, pp. 647–651, 2023, doi: 10.48175/IJARSCT-12098.

[3] S. Vitla, "Leveraging AI-Enhanced IAM, PAM, and IGA for Cybersecurity: A Cross-Industry Approach to Reducing Cyber Attacks," *J. Inf. Syst. Eng. Manag.*, vol. 10, pp. 70–91, 2025, doi: 10.52783/jisem.v10i16s.2568.

[4] S. Singamsetty and S. Singamsetty, "Hy-Search: A Hybrid Retrieval-Augmented Framework for Factual and Context-Aware Enterprise Knowledge Discovery," 2025. doi: https://doi.org/10.2991/978-94-6463-978-0_37.

[5] S. D. Rajesh Lingam, "Cost-Aware and Scalable Approaches for Large-Scale Model Evaluation in Enterprise Systems," *Milestone Trans. Artif. Intell.*, vol. 1, no. 1, pp. 119–137, March, 2026.

[6] E. Yılmaz, "Machine Identity Management in Modern Enterprise Security: Concepts, Challenges, and the Role of Privileged Access Management Systems," *Eng. Technol. Appl. Sci. Res.*, vol. 16, pp. 32369–32376, 2026, doi: 10.48084/etasr.16202.

[7] P. Schummer, A. del Rio, J. Serrano, D. Jimenez, G. Sánchez, and Á. Llorente, "Machine Learning-Based Network Anomaly Detection: Design, Implementation, and Evaluation," *AI*, vol. 5, no. 4, pp. 2967–2983, Dec. 2024, doi: 10.3390/ai5040143.

[8] S. Narang and A. Gogineni, "Zero-Trust Security in Intrusion Detection Networks: An AI-Powered Threat Detection in Cloud Environment," *Int. J. Sci. Res. Mod. Technol.*, vol. 4, no. 5, pp. 60–70, Jun. 2025, doi: 10.38124/ijsrmt.v4i5.542.

[9] M. Mehmood, R. Amin, M. M. A. Muslam, J. Xie, and H. Aldabbas, "Privilege Escalation Attack Detection and Mitigation in Cloud Using Machine Learning," *IEEE Access*, vol. 11, pp. 46561–46576, 2023, doi: 10.1109/ACCESS.2023.3273895.

[10] N. Shone, T. N. Ngoc, V. D. Phai, and Q. Shi, "A deep learning approach to network intrusion detection," *IEEE Trans. Emerg. Top. Comput. Intell.*, vol. 2, no. 1, pp. 41–50, 2018.

[11] A. Nerella and J. W. Sajja, "Responsible AI in Enterprise Applications: Balancing Innovation and Compliance," 2023, *MA Healthcare Ltd*. doi: 10.52710/cfs.744.

[12] E. Yilmaz, "Machine Identity Management in Modern Enterprise Security: Concepts, Challenges, and the Role of Privileged Access Management Systems," *Eng. Technol. Appl. Sci. Res.*, vol. 16, no. 1, pp. 32369–32376, 2026.

[13] S. K. Udayakumar, H. Ragothaman, and K. M. Khare, "A Novel Dataset and a Hybrid Ensemble Approach for Anomaly Detection in Enterprise-Access-Logs for Anomaly Detection," in *2025 International Conference on Data Science, Agents & Artificial Intelligence (ICDSAAI)*, 2025, pp. 1–6. doi: 10.1109/ICDSAAI65575.2025.11011761.

[14] C. Vijitphu and N. Utakrit, "A Machine Learning for Anomaly and Data Breach Detection System Using Python," in *International Conference on Cybernetics and Innovations, ICCI 2025*, 2025. doi: 10.1109/ICCI64209.2025.10987404.

[15] H. Uchida, K. Tominaga, H. Itai, Y. Li, and Y. Nakatoh, "Multi-Parameter Log Anomaly Detection with an Unsupervised Learning Approach," in *Proceedings of the 1st International Symposium on Parallel Computing and Distributed Systems, PCDS 2024*, 2024. doi: 10.1109/PCDS61776.2024.10743635.

[16] B. Torky, I. Karamitsos, and T. Najar, "Anomaly Detection in Enterprise Payment Systems: An Ensemble Machine Learning Approach," in *Lecture Notes in Operations Research*, 2024, pp. 11–23. doi: 10.1007/978-3-031-61589-4_2.

[17] S. Yan, S. Wang, Z. Chen, X. Jiang, and X. Cao, "CSLog: Anomaly Detection for Syslog Based on Contrastive Self-Supervised Representation Learning," in *2023 24st Asia-Pacific Network Operations and Management Symposium (APNOMS)*, 2023, pp. 165–170.

[18] R. U. L. A. Yeruva, D. Singh, S. Suddala, and N. Bhatt, "Augmented Data Management: Enhancing Cache Performance and Cybersecurity through AI and Mobile Integration," *J. Comput. Mech. Manag.*, vol. 5, no. 3, pp. 280–293, May, 2026, doi: https://doi.org/10.57159/jcmm.5.3.26691.