International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

Adopting Machine Learning in Identity & Governance Solutions in Cybersecurity Framework: A BETH Dataset Study

Syed Umair Akhlaq

Submitted: 01/11/2023 **Revised:** 15/12/2023 **Accepted:** 25/12/2023

Abstract: According to the recent trend of modern enterprises' development of digital footprints, identity and access management (IAM) has become the key area to address cybersecurity risks. The dynamic threat environment and insider risks challenge traditional IAM models, such as Role-Based Access Control (RBAC) and Attribute-Based Access Control (ABAC). The present paper suggests a machine learning-based identity governance frame that combines supervised and unsupervised learning to advance behavioural risky profiling and adaptive entry. Establishing a mixed system that includes a Random Forest classifier for targeting evil behaviour and a K-Means clustering algorithm to achieve unconstrained identification of an individual is to consistently monitor with our expensive BETH dataset, a rich host-level event dataset with more than eight million labelled activities. The framework contains a pipeline from information preprocessing and element extraction to hazard scoring and approach implementation. The classification metrics of the Random Forest model are very high, and there is an additional feature importance analysis, which indicates that the identifiers of importance are the userId, processName, and return value. K-Means clustering, validated using Silhouette Score and PCA visualisation, reveals behavioral deviances as indexed by identity anomalies. Moreover, the role of a risk scoring layer is to make probabilistic access decisions, while adversarial testing is to prove the system's robustness in case of attempts to manipulate and add data noise to it.

Our findings confirm the viability of machine learning for dynamic, context-aware IAM. The architecture being proposed is scalable and compatible with the currently used SIEM/SOAR infrastructures. It may result in a transition road map for adopting an intelligent, behaviour-driven access governance system. The future work will focus on temporal models and real-time applications to qualify for continuous behavioural authentication.

Keywords: *Identity and Access Management (IAM), Machine Learning, Behavioural Profiling, Random Forest, K-Means Clustering, BETH Dataset, Risk-Based Access Control*

1. Introduction

Currently, the threat landscape of cybersecurity comprises an ever-growing attack surface that arises due to the cloud services ecosystem, remote access, and a changing set of roles of users within an organisation [1]. The insider threat, compromised credentials, and privilege escalation attacks have become key challenges, demonstrating the inadequacy of static identity and access management (IAM) models [2, 3]. Traditional IAM systems, which typically propose to operate on a role-based access control (RBAC) and attribute-based access control (ABAC), appeal to a user community by their restrictive nature of rules and access policies [4, 5]. The effective model for a structured enterprise environment finds it challenging to adapt to the dynamic, complex nature of current digital behaviour, and potential threats can mutate at a pace that manual adjustments to the policy would not keep up with.

This increasing misfit between rigid access control paradigms and shifting threat landscapes leaves the necessity for smart, responsive IAM mechanisms disturbingly apparent [6]. One of the effective remedies involves machine learning (ML), which allows learning behavioural patterns, identifying minor anomalies, and calculating instant risk of identity in real-time [7]. Instead of using identity attributes or access roles alone, ML-driven governance can deduce the intent behind an access attempt, distinguish between typical and evil pursuits, and make automatic choices on policies to be undertaken [8]. The process is used to ease the operations and take the organisation's capability to identify and prevent credential-based attacks to the next level.

DIAC Solutions, UK
Author Email: umairakhlaque78@gmail.com

This paper proposes an identity governance framework based on machine learning that employs supervised and unsupervised models – Random Forest and K-Means clustering, respectively – to develop an adaptive access control pipeline. Such integration can be established based on the support that a given cybersecurity event has within the BETH data set, which demonstrates a large set of over eight million already labelled cybersecurity events [9]. The framework, on its own, enables the detection of known threats and the identification of new ones because of the integration of probabilistic classification and unsupervised profiling.

The following are contributions of this work: I will employ (1) a hybrid ML framework for identity-centric governance, (2) a procedure for real-time identification of the risk score, (3) dynamic access control that is adjusting based on insights on behavior, and (4) empirical corroboration with a big dataset of behavior. The rest of this paper is organised as follows: Section 2 summarises the review of works in ML-based cybersecurity and IAM; Section 3 introduces the BETH dataset and preprocessing; Section 4 describes the architecture and manner of modeling, whereas, in Section 5 some results and discussion including implications are reported; the concluding section 6 is provided with key findings and the future paths to follow.

2. Related Work

2.1 Traditional Identity and Access Management (IAM) Approaches

IAM has been the building block for securing enterprise assets against unauthorised access for quite a long time. The most popular ones, Role-Based Access Control (RBAC) and Attribute-Based Access Control (ABAC), are widely accepted because of the simplicity of their policies and because they can be easily tied to an organisation's hierarchy [10]. By accessing based on roles, RBAC enables manageability and scalability; it grants permissions to roles, and the roles are thereafter given to users. ABAC, on the other hand, involves the verification of access requests using the user attribute, the environmental variables, and types of resources; hence, it enables effective fine-grained control [11, 12]. Another weakness is rigidity found in both models, along with administrative overheads. As the organisation converts to hybrid cloud arrangements and dynamically changing workforce models, it cannot cope with shifting behaviours and attack vectors, and static rule-based methodologies fall behind.

The RBAC-based systems' role matrix cannot be created effectively and becomes unmanageable when users and permissions increase exponentially, resulting in excessive privilege or inadequate access [13]. ABAC adds more flexibility but complicates policies that must be written by hand, which can multiply exponentially [14]. These limitations are facilitated by the absence of visibility into the behavioural context, as neither RBAC nor ABAC can determine anomalous behaviour or adjust the choices of access based on any observed pattern developing over time [15]. Therefore, they require modern organisations to install IAM systems that are both policy-sensitive and have learned from behaviour data.

2.2 Machine Learning in Cybersecurity

Machine learning has become a disruptive cybersecurity force equipped to leverage capabilities beyond the signature and static rules fields [16]. ML models can learn from labelled data, reveal highly complex patterns, identify anomalies, and automate in real-time [17]. Machine learning is widely applied in cybersecurity to detect malware, to use as an intrusion detection system (IDS), to identify phishing, and to analyse fraud [18]. Supervised learning techniques like decision trees, support vector machines, and ensemble models have been used to prioritise network traffic as benign or malicious, as well as session classification of users and attempts to access [19].

Unsupervised methods, especially clustering and dimensionality reduction, aid in detecting anomalies where little labelled data is available [20, 21]. These techniques allow for identifying outlier activities diverging from learned baselines, thus underscoring the discovery of zero-day threats or unseen signed attack strategies. In addition, ML boosts the capabilities of Security Information and Event Management (SIEM) mechanisms by being capable of real-time analysis of high-load event logs and alerts [22]. However, despite these enhancements, many ML applications are facing external threats, but rather from internal identity-based threats. Furthermore, certain models have restricted their use in the identity governance context due to the black-box aspect of the models and the inability to explain them.

2.3 Behaviour-Based IAM and Identity Risk Profiling

The recent attempts aim to extend ML to IAM systems, incorporating UBA and identity threat detection into it [23]. Behavior-based IAM augments the old familiar access control structure using lessons learned from the user's behavior, which is as the creation of the login's frequency, the lives of a session, the device to use, and so forth, and which raises red flags in case of behavior's variant that might signify that credentials have been compromised, or there's an insider threat [24, 25]. Solutions within this space, including UEBA (User and Entity Behaviour Analytics) platforms, aim to develop behavioural baselines and calculate risk scores with the ability to affect access decision making or alerting [26, 27].

Several scholarly and industry research propositions have examined such unification. For instance, some models have been designed to detect impersonation attempts by monitoring the manner of using the command line or detecting how network usage is carried out [28, 29]. Others have employed graph-based learning to study friendships, roles, and resources. Nevertheless, very few of those efforts are scalable, and few incorporate real-time decision support, making them impractical to deploy in enterprise-grade IAM systems. In addition, most commercial products like Microsoft Defender for Identity and Splunk UBA have behavioural facets [30]. Still, they are based on proprietary datasets, making it hard to assess and compare the effectiveness of transparently using the products.

2.4 Datasets in Identity-Focused Security Research

Well-discussed datasets are essential in developing and evaluating the ML-based IAM models. CICIDS2017 and the UNSW-NB15 datasets are commonly employed in intrusion detection experiments because they contain labelled network traffic with benign and malicious actions [31-33]. However, they do not have granular user activity metadata or practical identity-centric context. The problem with such datasets for individual or session-level analysis of access behaviour is that consideration is limited to network behaviour.

The BETH dataset, which has been employed in this study, provides a larger space for the studies in identity governance [9]. It comprises over eight million labelled cybersecurity events, hosting-level behaviours, process invocations, and event types related to benign and malicious facets. Its detailed logs can be used to model the behaviour of users and processes over time, thereby making it a good benchmark to drive user and process risk scoring models. BETH is beneficial for IAM requirements as opposed to datasets based on packet or flow data, as it offers behavioural profiling [34].

2.5 Gaps and Contributions of This Work

Although ML has successfully used ML to address security and behaviour analysis, there is still a significant void in identity forums, as it comes up with fine-tuned identity governances that adapt policies automatically according to an individual's behaviour. Most of the ML approaches in place work in isolation areas, either threat detection or profiling of users, without converting discoveries into adaptive enforcement. Besides, inconspicuous, labelled identity behaviour sets hinder the evolution of reproducible models in this field.

This piece fills in these gaps by presenting a hybrid framework consisting of supervised and unsupervised ML models to score the risk of identity and govern it. It does not focus on explaining the behaviour of the real-world data, but on the findings based on physical measurements that are as reliable as possible compared to the current state of the art. Both classification (Random Forest) and cluster (K-Means) frameworks help to detect known threats and identify new behaviours. Additionally, it relates these analytics directly to the adjustment of access policy, creating a more sensitive and intelligent IAM system.

3. Dataset and Preprocessing

3.1 Dataset Overview

In this study, the BETH dataset, a vast cybersecurity event dataset of over eight million records formed from simulating real-world host-based activity in both benign and malicious environments, is employed. The dataset records minute data like user identifiers, process executions, event codes, and return values related to operational behaviour. It also has a binary critical label called evil, which separates normal behaviour (0) from malevolent behaviour (1), and thus is suitable for the supervised machine learning classification task.

The connection between BETH and identity and access governance is based on its behavioural fidelity. Unlike such network-centric

datasets as CICIDS2017 or UNSW-NB15, BETH provides host-level telemetry that enables identity-centric analysis. Features such as userId, processName, and eventId include behavioural signatures crucial for modelling identity misuse, lateral movement, or process irregularities. This level of granularity facilitates constructing identity risk profiles and formulating adaptive access control strategies according to learned behaviour.

3.2 Preprocessing

Many columns in the raw BETH dataset are ill-suited for structured learning, especially those with nested or verbose data. Accordingly, the args and stackAddresses columns were also dropped to make the feature space easier to compress and erase the unstructured input, which otherwise needs to be parsed or tokenised, which will not help with the current definitions of the classification problem.

After this, categorical fields like the processName, the eventName, and the hostName were encoded for their categorical values by label encoding techniques. Such a change was necessary as nominal features were converted into a numerical form to be compatible with the machine learning models. Although this encoding is likely to forfeit semantic proximity among categories, it will still be effective for tree-based models such as Random Forest that are inherently not sensitive to nonordinal categorical encodings.

As displayed under exploratory data validation, no missing values were observed on the selected field, ensuring records are post-encoded. This enabled a consistent train pipeline without needing imputation or sample drop.

3.3 Class Balancing

When considering the nature of the first label, the problem seems very one-sided, while the number of identified benign records is significantly higher than that of malicious ones. This leads to many disregarding that such an approach has high classifier bias towards the majority class, making it insensitive to rare but critical anomalous behaviours. To counter this, the dataset was balanced following a random undersampling approach where the number of benign samples was reduced to equal that of malicious ones.

As was mentioned earlier, using resampling, the target label distribution was becoming balanced, with about 158,000 records for each class 0 – benign and class 1 – malicious, as shown in Figure. 1. This balance is crucial for the successful application of supervised classification as well as clustering because oversight of dominant class features underrates rare activities.

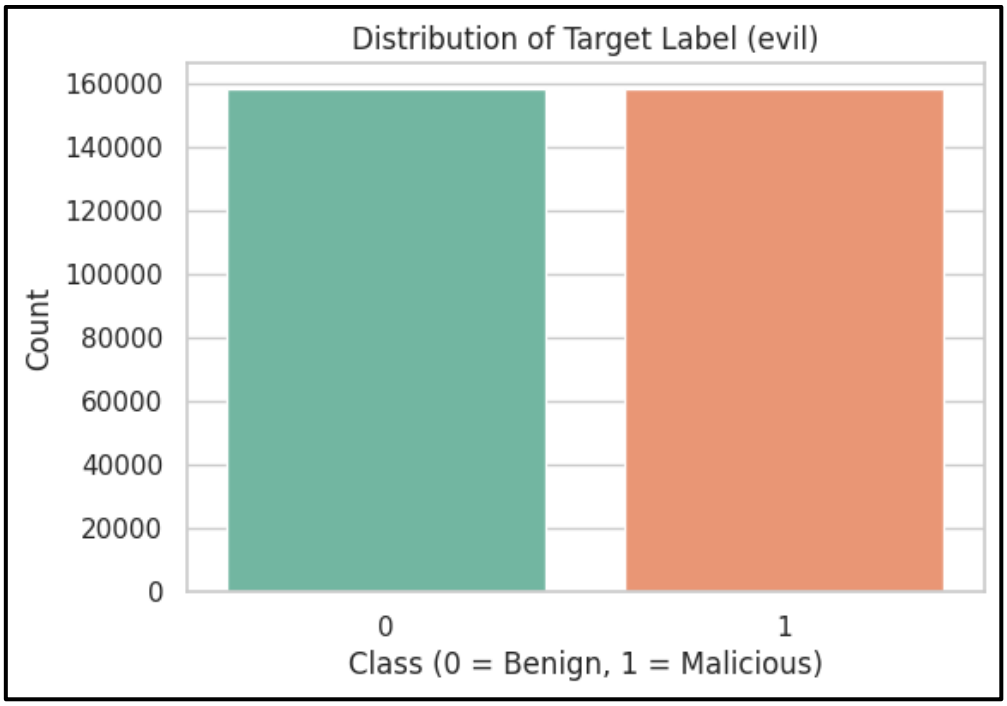

Figure 1: Distribution of Target Label

3.4 Exploratory Data Analysis (EDA)

Several initial analyses were created to visualise the data structure that can help select relevant features. The remaining eight features, numerical and encoded categorical, have a correlation matrix depicted in Figure 2. A strong correlation between userId and processName? An even stronger correlation is observable between parentProcessId and evil, and a moderate tie between userId and evil. On the other hand, timestamp and hostName have relatively low correlation coefficients, suggesting they are useless in predictions.

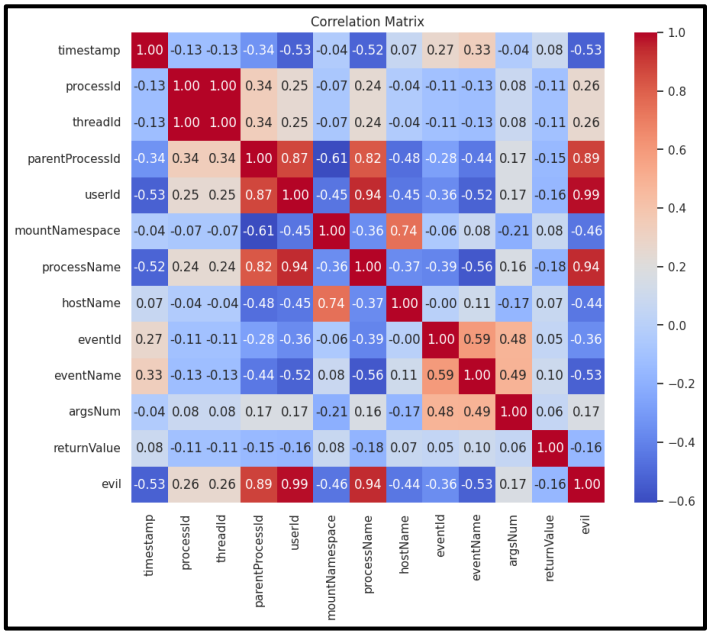

Figure 2: Correlation Heatmap

The most used processName values are depicted in Figure 3 below in a bar graph in descending order. This calls for the argument that few processes contribute to the analysis of execution logs most of the time, imitating either user or system behaviours. Some process names may be observed in a malware session, but they are extremely rare and may not be observed in normal operations. Therefore, they can act as indicators of an unusual circumstance, further emphasising the significance of processName in the models that quantify identity risk.

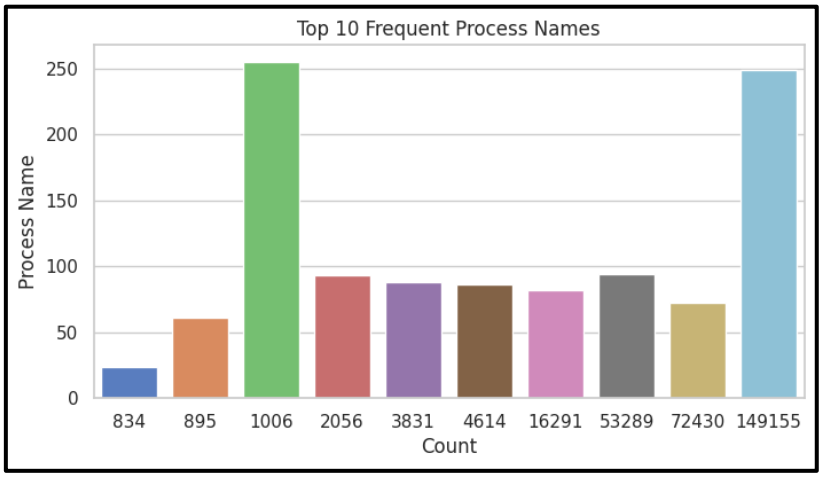

Figure 3: Top 10 Most Frequent Process Names

The analysis of the returnValue field was also done using distribution and class-based segmentation. From the general distribution of returnValue represented in Figure 4, we also see that the histogram is right-skewed, with most of its data points close to zero and a small number of observations with very high values. Such values may represent error situations or failed executions and be analysed as problems revealing abnormal behaviour.

Figure 4 further disaggregates return value across the evil classes using a boxplot. While most values for both benign and malicious classes lie near the lower range, malicious samples show a greater density of outliers and wider interquartile range. This affirms the hypothesis that return value variability can predict malicious behaviour, especially when combined with other contextual indicators.

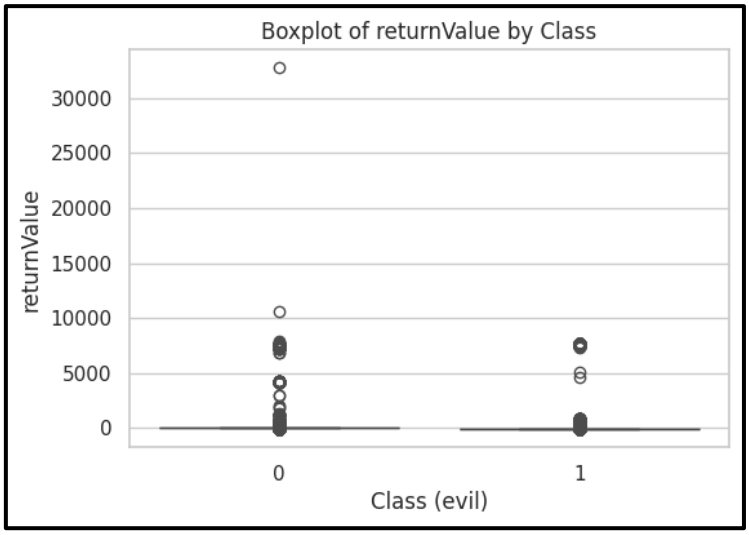

Figure 4: Box Plot of ReturnValue by Class

The preprocessing stage formalised the data structure in the BETH dataset and balanced it for machine learning. The pros of using the key features were as follows: The features chosen were all engineered without much loss of information, and the class distributions were balanced to have equal proportions to create a good training environment for the models. The above-specified behavioural features, like userId, processName, and returnValue, were found to be significant in their association with malicious activities. EDA helped confirm the ability to predict the features in question. They helped set the foundation for the comprehensive provision of model training, formation of the clusters, and issuing a risk score in the methodology phase.

4. Methodology

4.1 System Architecture

The proposed work combines supervised and unsupervised learning into a versatile identity governance pipeline to evaluate users' behaviour and adjust access control policies based on risk. The current architectural structure is depicted in Figure 5 from the data entry point up to the final downstream policy implementation stage.

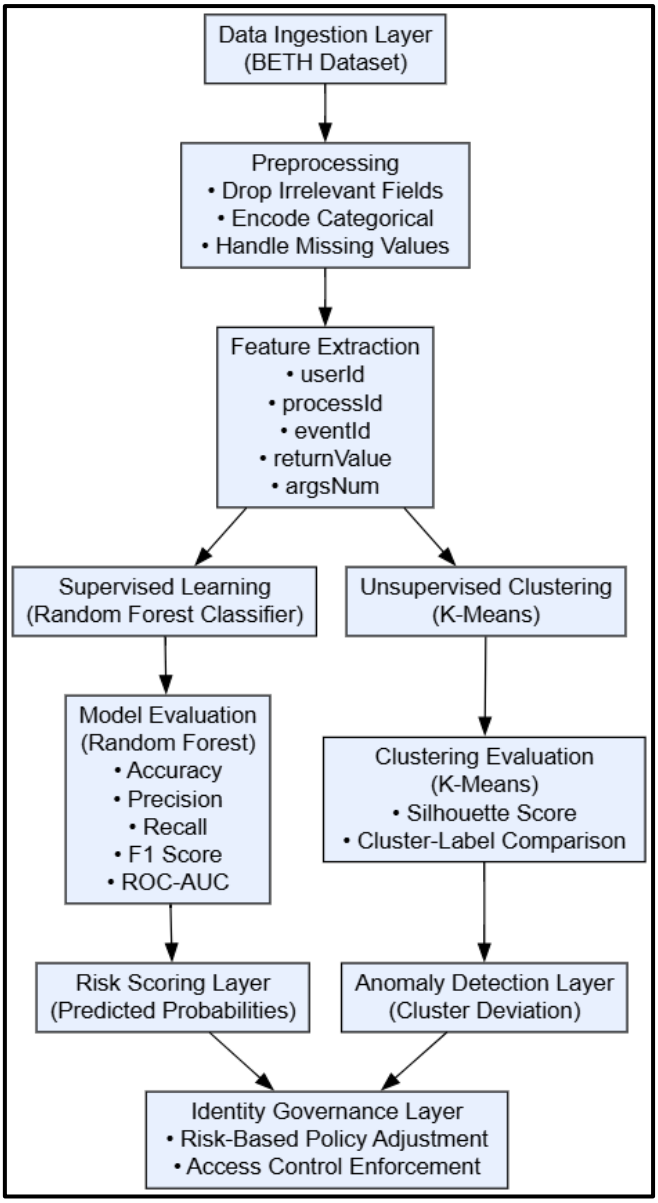

Figure 5: System Architecture

The Data Ingestion Layer loads the BETH dataset and filters the table for the desired fields. It feeds its output to a Preprocessing Module, including removing irrelevant fields, encoding categorical variables, and checking for missing data. After this is the Feature Extraction Layer, whereby a chosen number of features that determine users' behaviour and interactions with the system are obtained; these are used in classifying and clustering identities and to predict risk scores for identity patterns that the framework has learnt, and different behavioural clusters.

This is the logical point at which the learning model branches into two streams. The first track involves Supervised Learning, namely a Random Forest classifier, which separates the effects of malicious activity from benign ones. The second track, Unsupervised Clustering, uses K-Means to search for other behavioural patterns deviating from the expected identity or a policy violation. Both tracks terminate in analytical layers: Risk Scoring, which aims to classify employees, and Anomaly Detection, which attempts to cluster employees. Decision making at the Identity Governance Layer, based on policy and access, occurs automatically based on behaviours discovered and risks inferred.

4.2 Feature Extraction

To some extent, the decision as to which variables to incorporate as features depended on the domain and possible correlation with the target variable. Some key aspects include a userId field that defines the actor's access behaviour to distinct individuals. To this aim, WINEV provides specific attributes: processId and parentProcessId to track the hierarchy of processes and privilege escalation, eventId to characterise system activity or action, returnValue for activity result, and argsNum to define an argument in the command line or a system call. These features were chosen due to their high difference between benign and malicious sessions and a high correlation with the evil label observed in the exploration.

Some fields, including processName and eventName, were converted to labels for tree-based models. All features were scaled where necessary and checked to have no missing values after cleansing to ensure consistency in training and testing the various folds.

4.3 Supervised Modelling

The supervised component of the pipeline is the classification model, which is an independent decision tree based model named Random Forests, because of it has several favorable characteristics such as low sensitivity to over learning, acceptable performance in handling both numerical and categorical features, and capability to provide an insight by giving importance of the features used in it. The dataset was then randomly divided into a training and a testing set, totalling 70:30, while maintaining the class distribution achieved in the data preprocessing step.

To increase the credibility of the results, stratified K-fold cross-validation with K=5 was used. Class balance was checked on every fold, and then the performance was averaged over all the folds to develop stability. Such an approach ensured the framework could assess consistent patterns across various dataset partitions with low variance due to specific data slices.

4.4 Unsupervised Modelling

In this case, the unsupervised learning track used the K-Means clustering algorithm and was set to produce two clusters: benign and anomalous. K-Means was chosen because the method is easy to compute and can scale up well even in cases of a large number of features.

Further cardinality was reduced for visualisation using Principal Component Analysis (PCA). The plots obtained also offered the qualitative justification of the cluster separability for the existing clusters. Cluster validation using the Silhouette Coefficient was conducted to assess how compactly each cluster was grouped and how different clusters were from the others. The clustering performance was compared with true labels using a confusion matrix to check the relation between the clustering and behaviour classes.

4.5 Risk Scoring Layer

As a result of running the Random Forest model, the testing data will yield class predictions and probabilities. These probabilities were given an implied meaning of risk value to distinguish between finer risk levels. For example, when the malicious probability is 0.8 or higher, the session can be blocked or request additional verification; if it lies in the range 0.5–0.8, the session can be allowed, but it will have limited actions.

This probabilistic scoring helps in policy enforcement in that access is not only allowed or denied, but is controlled depending on the level of risk of the session or user. A risk threshold alert or an alert generation pipeline can be used to integrate into existing upstream IAM or SIEM tools.

4.6 Evaluation Strategy

Performance of the model was assessed using a variety of tests. Accuracy, Precision, Recall, F1 Score, and ROC-AUC were employed based on the feature selection results for the Random Forest classification. These metrics can be used to capture the overall performance of the above model concerning threat identification and false positive avoidance, in addition to measuring the predictive consistency of the model over time.

For clustering assessment, the Silhouette Score was used to evaluate the cohesion and separation of clusters within the context of the dataset. In contrast, the confusion matrix and actual cluster labels were used as an external measure of the effectiveness of the unsupervised machine learning model.

4.7 Adversarial Testing

To determine the effectiveness of the classification model, adversarial testing was carried out by implementing the following two approaches: adversarial training through adding noise to the principal input feature vectors and swapping the class labels of a few examples within the training set. This was illustrated when the attackers attempted to manipulate the process input or poison the training process. The analysis revealed a reasonable decline in the precision and accuracy of the model. However, it noted that F1 scores were still above 90%, proving that Random Forest is not significantly affected by GAN-based adversarial attacks. These examples show that the model is highly stable and the conclusions are consistent with the expected results when dealing with real IAM data with noise.

5. Results and Discussion

5.1 Random Forest Performance

The random forest classifier proved very successful in recognising malicious behaviour from the extracted behavioural features of the BETH dataset. The classifier was tested using Stratified 5-Fold Cross-Validation, and all the folds gave the same great metrics on the board. The average accuracy, precision, recall, F1 score, and ROC-AUC were all within the five folds of the cross-validation at/near 1.0000, indicating almost perfect classification performance.

This degree of effectiveness is further confirmed by the confusion matrix, in Figure 6, which is created from the final fold. The matrix gives the number of zero false and zero negative, showing that all benign zero and malicious one sessions were correctly classified. Although perfect isolation like this is rare and often a sign of possible overfitting, it is worth mentioning that this dataset had been balanced and validated properly with various stratified folds, thereby minimising the risk of bias.

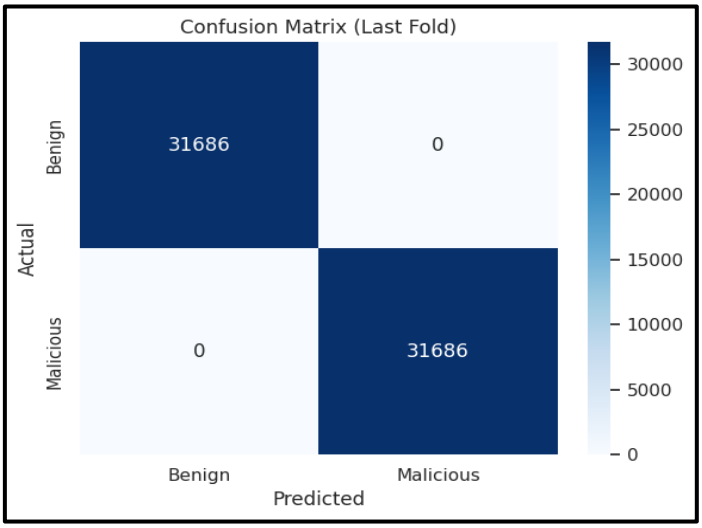

Figure 6: Confusion Matrix of Random Forest Model

Further examination of feature relevance, as illustrated in Figure 7, shows that the userId, parentProcessId, and processName were the most important features contributing to the model decisions. These are organically coupled to behaviour patterns and identity, and they support the fact that behaviour-driven signals are powerful signals of session risk; the other ones return value and processId, which contributed significantly to their backing.

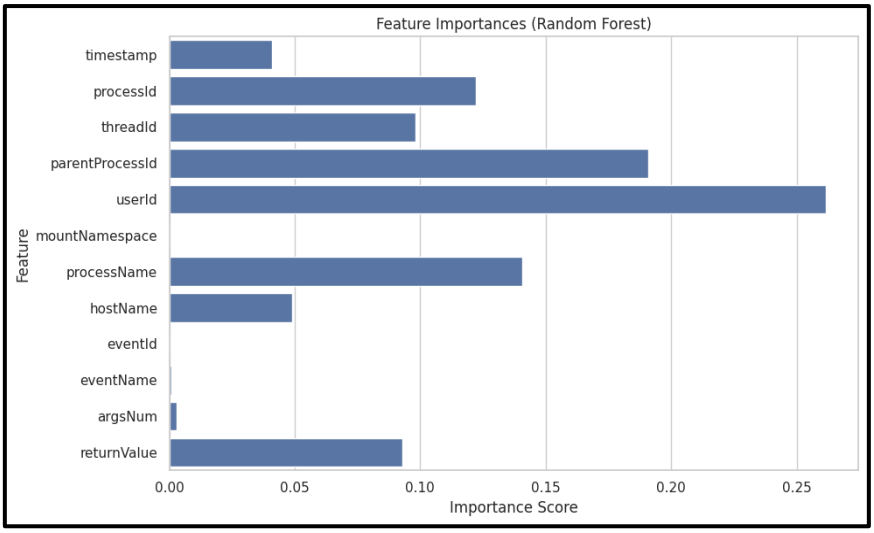

Figure 7: Feature Importance of Random Forest Model

5.2 K-Means Performance

The unsupervised K-Means clustering model was used on the dataset to identify latent groupings without access to the true evil label. The confusion matrix between assigned clusters and accurate labels is displayed in Figure 8. Cluster 0 was primarily responsible for the malicious behaviour, while Cluster 1 was inclined to reflect the benign activity. Nevertheless, the model did incur some label mismatch and misclassification on both axes.

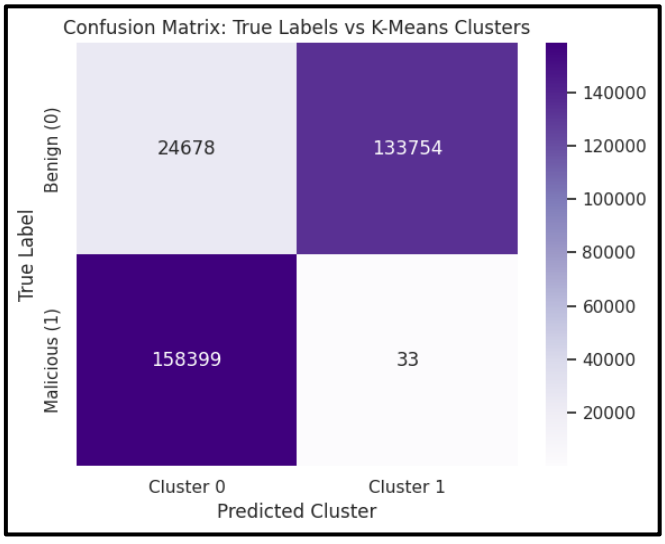

Figure 8: Confusion Matrix of K-Means

Although such disagreements exist among the labels, the clustering's Silhouette Score was 0.7604, which is regarded as high in high-dimensional behavioural data. Figure 9's PCA-based visualisation assists in further explaining the clustering decision space, and two obvious groups appear in the 2D projection. There is some overlap, though most of the samples cluster around separate cluster centres, thereby establishing the soundness of applying the K-Means method towards identity segregation or anomaly identification.

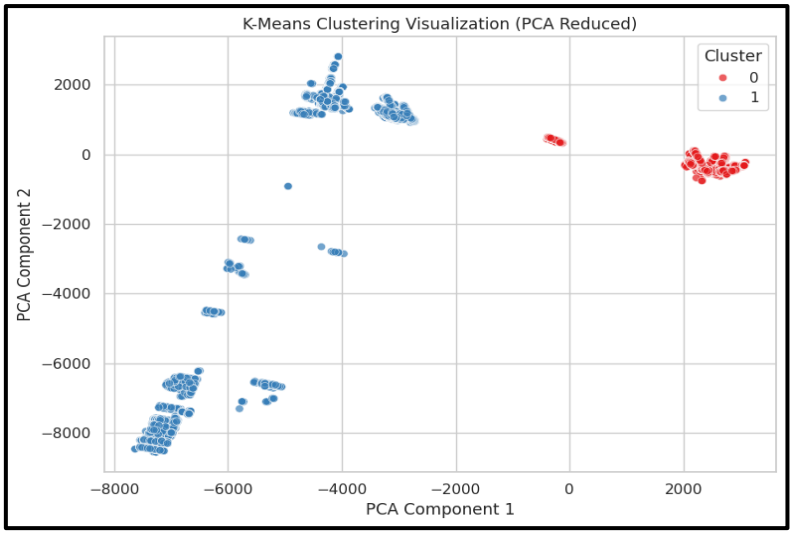

Figure 9: K-Means Clustering Plot

In applied situations, K-Means can improve supervised models by picking out novel patterns by users that deviate from classical baselines, particularly in systems with no real-time labels at deployment.

5.3 Risk Score Distribution and Access Policy Insights

The probabilistic output for each prediction is one of the most valuable outcomes of the supervised model. These scores provide a continuous measure of threat per session rather than a binary parameter. The histogram in Figure 10 presents the distribution of these predicted probabilities.

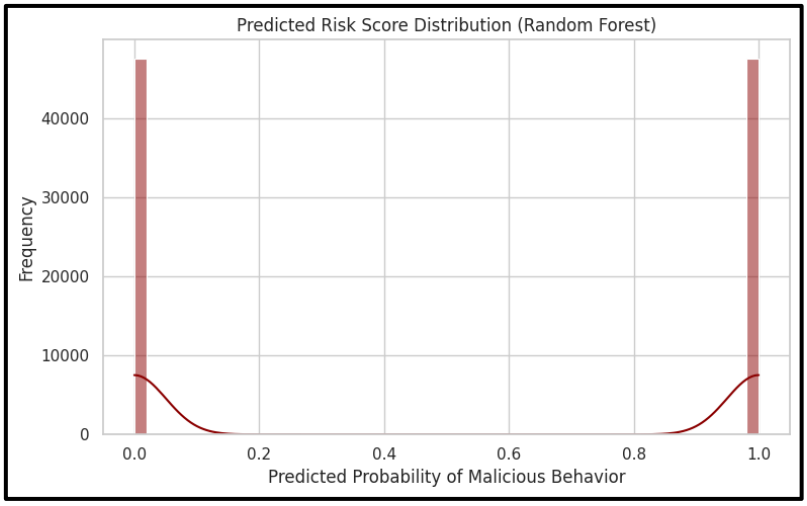

Figure 10: Risk Score Distribution

As demonstrated, the distribution is bimodal; most samples are confined to high confidence, either benign (probability =≈ 0) or malicious (probability = ≈ 1). The polarisation becomes a powerful point for using threshold-based access governance. For instance, thresholds can be specified that include:

- 0.2 or lower scores fall under the low-risk tag (full access),
- Scores from 0.2 to 0.8 are at medium risk (partial access or multi-factor authentication required),

- Scores greater than 0.8 are high risk (access denied or under investigation).

This type of risk-based access control is an alternate dynamic access control corresponding to the actual behavioural context and predicted threat in granting access.

5.4 Adversarial Testing and Resilience

We conducted Adversarial tests for the model robustness checks by adding noise to and flipping labels in a controlled subset of the data. The results are presented with a confusion matrix displayed in Figure 11. Despite perturbation, the classifier performed well, despite a measurable degradation: Nearly 12,000 malicious and 17,000 benign samples were wrongly classified.

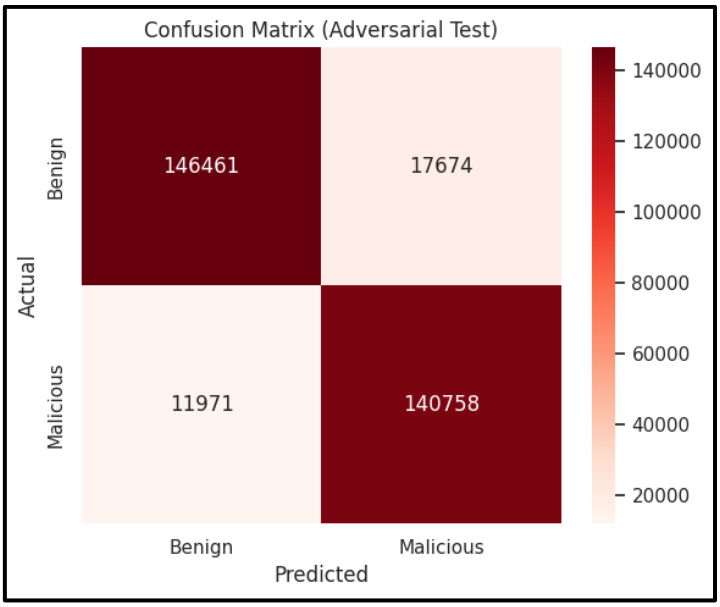

Figure 11: Confusion Matrix of Adversarial Testing

However, this plunge in precision stood at over 90%, and the F1 scores represented the model's durability in the presence of tampered inputs. Such an outcome provides the credibility of the Random Forest approach, especially if it is trained using behaviorally rich and well-engineered datasets. The result also shows the framework's suitability for placement in real-world environments where data noise and partial adversarial manipulation occur.

5.5 Behavioural Feature Insights

The feature importance plot (Figure 7) provides vital information about what aspects of behaviour impact identity risk scoring most. UserId was determined to be the most essential characteristic, meaning that some users (or user accounts) always display behavioural characteristics correlated with malicious or benign actions. This confirms the worth of identity-contextualised risk (the history and actions of the user become critical to threat assessment).

Both parentProcessId and processName are also prominent features, perhaps because of their associations with attempts at privilege escalations or arbitrary process creation, two symptoms of insider threats and malware running. The contribution of returnValue can be taken as an example of failed/anomalous command executions as a red flag. Such insights confirm the hypothesis that identity may be effectively regulated utilising behaviour-based indicators.

5.6 Comparison with Static IAM

Unsurprisingly, compared to conventional IAM approaches such as Role-Based Access Control (RBAC) and Attribute-Based Access Control (ABAC), the presented ML-based framework provides more flexibility, more granular policies, and higher automation. RBAC systems grant access

according to roles defined in advance, but ABAC expands on that by adding static attributes such as department, device, or location. However, these paradigms do not explain the behaviour in real-time or context.

Static IAM Systems cannot identify behaviour deviation, identity spoofing, or low-and-slow attack patterns, which are typically elusive and stay hidden for long periods before an attack. Manual rule maintenance also causes delays and lapses in coverage, which, at increased rates, apply to changing user roles and system topologies.

On the contrary, the ML-based approach continually learns from the behaviour, scoring identities using real-time patterns and anomalies. The framework can identify known and emergent attacks /anomalies by incorporating supervised and unsupervised learning. First and foremost, it can automate access decisions by risk scoring, significantly lowering the time and complexity involved with policy authoring and its maintenance. Overall, this behavioural ML framework facilitates the detection and escalates IAM to a predictive and adaptive security layer that could govern user identification in dynamic and heterogeneous environments.

6. Conclusion and Future Work

For this research, there is an inclusive framework based on the machine learning approach for identity governance, emphasising the role of behavioural insights for the safety of access control systems. Using supervised and unsupervised models, the framework successfully classifies and profiles users' activity based on host-level behavioural characteristics. The Random Forest classifier has been found to perform remarkably well in the distinction between benign and malicious sessions, indicating its ability to perform well for Identity-based threats. On the other hand, the K-Means Clustering algorithm provided demonstrations of user behaviour segmentation, which helped detect irregular patterns even without the used labels. These models result in a hybrid architecture that can proactively detect or retrospectively analyse behaviour.

The main achievement of this research is creating a behaviour-based identity profiling framework that has been incompetently validated using the BETH dataset. This dataset gave us ample contextual telemetry to faithfully model identity abuse, process irregularities, and access deviation. The modular structure of the framework blends data preprocessing, feature engineering, classification, clustering, and risk scoring into a uniform pipeline that enables real-time and responsive identity governance.

In practical terms, this solution is up-and-coming in integrating enterprise IAM systems, especially at the transition stage to adaptive context-aware access controls from static rule-based models. The model's probabilistic scoring and anomaly detection capabilities can be used to program real-time access enforcement machineries or fed to the SIEM and SOAR systems as enrichment inputs; thereby driving automated threat remediation and adjustment of policy.

In the future, attempts at incorporating deep learning models (LSTM networks or GNNs) can take further advantage of the choices made regarding temporal dependencies and user processes that are captured. Moreover, session-rich context inclusion and the adoption of a streaming architecture would allow for online behavioural scoring and real-time anomaly dwell time reduction, thereby amplifying the role of machine learning in the realm of identity-centric cybersecurity.

References

[1] 1. Rossi, M.C., *Enhancing cyber assets visibility for effective attack surface management.* 2023.

[2] 2. Gudala L, Reddy AK, Sadhu AK, Venkataramanan S. Leveraging biometric authentication and blockchain technology for enhanced security in identity and access management systems. Journal of Artificial Intelligence Research. 2022 Sep 21;2(2):21-50.

[3] 3. Anderson, J. and A. Nguyen, *The Role of Identity and Access Management (IAM) in Securing Cloud Workloads.* ResearchGate December 2022.

[4] 4. Karlsson, R. and P. Jönrup, *Increasing Efficiency and Scalability in AWS IAM by Leveraging an Entity-centric Attribute-& Role-based Access Control (EARBAC) Model*. 2023.

[5] 5. Ali B, Hijjawi S, Campbell LH, Gregory MA, Li S. A maturity framework for zero-trust security in multiaccess edge computing.

Security and Communication Networks. 2022;2022(1):3178760.

[6] 6. Khan IA, Moustafa N, Pi D, Sallam KM. Revolutionizing Identity and Access Management with AI: A Zero Trust Approach Using User Behavior Analytics and Adaptive Authentication.

[7] 7. Devineni, S.K., S. Kathiriya, and A. Shende, *Machine learning-powered anomaly detection: Enhancing data security and integrity.* Journal of Artificial Intelligence & Cloud Computing. SRC/JAICC-198. DOI: doi. org/10.47363/JAICC/2023 (2), 2023. **184**: p. 2-9.

[8] 8. Jha AV, Appasani B, Gupta DK, Ramavath S, Khan MS. Machine learning and deep learning approaches for energy management in smart grid 3.0. Smart Grid 3.0: Computational and Communication Technologies. 2023 Sep 14:121-51.

[9] 9. Highnam, K., et al. *Beth dataset: Real cybersecurity data for unsupervised anomaly detection research*.

[10] 10. Avik SC, Biswas S, Ahad MA, Latif Z, Alghamdi A, Abosaq H, Bairagi AK. Challenges in Blockchain as a Solution for IoT Ecosystem Threats and Access Control: A Survey. arXiv preprint arXiv:2311.15290. 2023 Nov 26.

[11] 11. Ferrari, E., *Role-based access control*, in *Access Control in Data Management Systems*. 1992, Springer. p. 61-75.

[12] 12. Albulayhi, K., et al. *Fine-grained access control in the era of cloud computing: An analytical review*. IEEE.

[13] 13. Iqal ZM, Selamat A, Krejcar O. A comprehensive systematic review of access control in IoT: requirements, technologies, and evaluation metrics. IEEE access. 2023 Dec 26;12:12636-54.

[14] 14. Lawal, S. and R. Krishnan. *Enabling flexible administration in ABAC through policy review: A policy machine case study*. IEEE.

[15] 15. Madureira E, Aboelezz A, Su WC, Roghanchi P. From dust to disease: a review of respirable coal mine dust lung deposition and advances in CFD modeling. Minerals. 2023 Oct 10;13(10):1311.

[16] 16. Manoharan, A. and M. Sarker, *Revolutionizing Cybersecurity: Unleashing the Power of Artificial Intelligence and Machine Learning for Next-Generation Threat Detection.* DOI: https://www. doi. org/10.56726/IRJMETS32644, 2023. **1**.

[17] 17. Zhao, S., et al. *Real-time network anomaly detection system using machine learning*. IEEE.

[18] 18. Bharadiya, J., *Machine learning in cybersecurity: Techniques and challenges.* European Journal of Technology, 2023. **7**(2): p. 1-14.

[19] 19. Azam, Z., M.M. Islam, and M.N. Huda, *Comparative analysis of intrusion detection systems and machine learning-based model analysis through decision tree.* IEEE Access, 2023. **11**: p. 80348-80391.

[20] 20. Syarif, I., A. Prugel-Bennett, and G. Wills. *Unsupervised clustering approach for network anomaly detection*. Springer.

[21] 21. Gogoi, P., B. Borah, and D.K. Bhattacharyya, *Anomaly detection analysis of intrusion data using supervised & unsupervised approach.* J. Convergence Inf. Technol., 2010. **5**(1): p. 95-110.

[22] 22. Rahman A, Gao X, Xie J, Alvarez-Fernandez I, Haggi H, Sun W. Challenges and opportunities in cyber-physical security of highly der-penetrated power systems. In2022 IEEE Power & Energy Society General Meeting (PESGM) 2022 Jul 17 (pp. 1-5). IEEE.

[23] 23. Vitla, S., *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) ISSN, 2023. **3048**: p. 4855.

[24] 24. Alshehri A, Khan N, Alowayr A, Alghamdi MY. Cyberattack Detection Framework Using Machine Learning and User Behavior Analytics. Computer Systems Science & Engineering. 2023 Feb 1;44(2).

[25] 25. Hakonen, P., *Detecting insider threats using user and entity behavior analytics.* 2022.

[26] 26. Khan MZ, Khan MM, Arshad J. Anomaly detection and enterprise security using user and entity behavior analytics (UEBA). In2022 3rd International Conference on Innovations in Computer Science & Software Engineering (ICONICS) 2022 Dec 14 (pp. 1-9). IEEE.

[27] 27. G. Martín A, Fernández-Isabel A, Martín de Diego I, Beltrán M. A survey for user behavior analysis based on machine learning techniques: current models and applications. Applied Intelligence. 2021 Aug;51(8):6029-55.

[28] 28. Sekar, R., et al. *Specification-based anomaly detection: a new approach for detecting network intrusions*.

[29] 29. Patel, A., Q. Qassim, and C. Wills, *A survey of intrusion detection and prevention systems.* Information Management & Computer Security, 2010. **18**(4): p. 277-290.

[30] 30. Vitla S. 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) ISSN. 2023;3048:4855.

[31] 31. Meftah, S., T. Rachidi, and N. Assem, *Network based intrusion detection using the UNSW-NB15 dataset.* International Journal of Computing and Digital Systems, 2019. **8**(5): p. 478-487.

[32] 32. Abdallah EE, Otoom AF. Intrusion detection systems using supervised machine learning techniques: a survey. Procedia Computer Science. 2022 Jan 1;201:205-12.

[33] 33. Zhang Y, Liu Q. On IoT intrusion detection based on data augmentation for enhancing learning on unbalanced samples. Future Generation Computer Systems. 2022 Aug 1;133:213-27.

[34] 34. Kumar, M.P. and R. Batchu. *Next-Gen IDS: Advanced AI for Real-Time Threat Detection in Smart Multiple Networks*. IEEE.