Suman Kumar Sanjeev Prasanna[1*], Xiaojun Ruan[2]

A Unified Hybrid Machine Learning Architecture for Robust Identity Anomaly Detection in Large-Scale Digital Ecosystems

Journal of Electrical Systems

Abstract: Identity anomaly detection in large-scale digital ecosystems is challenged by non-stationary behavioral dynamics, high-dimensional heterogeneous feature spaces, and limited availability of labeled anomaly instances. Existing approaches rely on isolated supervised or unsupervised models, restricting their ability to simultaneously detect known attack signatures and emerging behavioral deviations.This work introduces a unified hybrid learning architecture that addresses these limitations by integrating supervised classification, reconstruction-based unsupervised modeling, and temporal representation learning within a single optimization framework. By jointly optimizing discriminative and generative objectives through sliding-window aggregation and ensemble decision fusion, the architecture captures both short-term behavioral fluctuations and long-term identity evolution patterns. Empirical evaluation on large-scale identity interaction datasets demonstrates that the proposed framework achieves 97.3% accuracy and a 97.1% anomaly detection rate, outperforming strong supervised and unsupervised baselines by up to 7.7% in ADR. These results indicate that multi-paradigm temporal representation learning substantially enhances robustness to previously unseen anomalies under non-stationary conditions, providing a scalable foundation for identity-centric anomaly detection in complex digital environments.

Keywords: Hybrid Machine Learning; Identity Anomaly Detection; Temporal Representation Learning; Supervised–Unsupervised Integration; Behavioral Modeling; Ensemble Learning; Non-Stationary Data; Reconstruction-Based Modeling; Digital Identity Security; Large-Scale Anomaly Detection

1. Introduction

Digital ecosystems increasingly rely on complex identity infrastructures that span cloud services, social platforms, financial systems, healthcare networks, and interconnected IoT environments [1]. In such systems, digital identity is no longer a static entity but a constantly evolving behavioral construct defined by user interactions, access patterns, contextual signals, and cross-platform activity [2]. These changes have greatly heightened exposure to identity anomalies such as impersonation, synthetic identities, coordinated fraudulent behaviors, and stealthy account manipulation [3]. Traditional identity security mechanisms-essentially based on rule engines, static authentication credentials, and predefined threat signatures-are inherently limited in the detection capability for adaptive, data-driven, and evolving patterns of anomalies [4]. These limitations become particularly critical in large-scale digital ecosystems, where high velocity, heterogeneity, and interconnectivity create the conditions whereby anomalous identity behaviors may lie hidden within the noise of normal operations [5]. Hence, there has been a rise in the need to have intelligent detection systems that are capable of learning identity dynamics directly from the data, and also able to identify normal behavior profiles and deviations [6]. Machine learning gives this potential to make the necessary shift by allowing systems to model non-linear relationships and dependencies in identity integrity [7]. Of course, many existing solutions are disjoint, domain-dependent, and/or limited to theoretical validation, and yet it is not clear how these solutions can be made scalable in the real digital world [8].

The paper positions itself as a results-oriented contribution to overcome such limitations by designing and empirically evaluating machine learning architectures that are optimally tailored for identity anomaly detection. In contrast to pure conceptual modeling, this study focuses on experimental validation with structured digital identity data, showing quantifiable improvements in anomaly detection performance regarding accuracy, robustness, and generalization. These architectures can be applied to heterogeneous identity features, thus

[1]Department of Computer Science, California State University, East Bay, Hayward, USA
ssanjeevprasanna@horizon.csueastbay.edu
[2]Department of Computer Science, California State University, East Bay, Hayward, USA
xiaojun.raun@csueastbay.edu

achieving consistent detection performance in complex digital environments. Supervised and unsupervised learning paradigms, combined with ensemble learning within one framework, avoid model evaluations in isolation and offer comparative performance insights for several learning strategies[9]. The results evidence the viability of such hybrid architectures, wherein various identified, as well as unknown, anomaly patterns can be effectively captured, thereby tackling an important problem that is inherent in any real-world identity frameworks, in which threat behaviors are always changing [10]. Notably, this work demonstrates an important contribution, that of verifying the viability of data-driven identity anomaly detection, within a digital ecosystem [11]. The obtained results form a basis for future smart identity security infrastructures that are adaptive, robust, and responsive to new threats, thus placing machine learning, rather than as an abstract technology, into a practical application for digital identity security [12]-[13].

The paper aims to create advanced architectures in machine learning for the detection of identity anomalies in digital ecosystems. The research applies to various heterogeneous digital identity ecosystems, such as multi-platform interaction systems, cloud computing, data-centric digital infrastructures, etc., where identity behaviors are dynamic, distributed, and constantly changing. Rather than focusing solely on identity verification mechanisms, the research emphasizes the concept of identity as a behavioral/contextual process to enable the detection of anomalies through data-centric pattern learning. The framework is best suited to support large scales of identity data while maintaining detection accuracy and reliability. By focusing on the empirical evaluation of the research, the authors are approaching an application-specific research paper that closes the gap between theoretical models of machine learning and actual digital identity security systems.

The inspiration to undertake the research is derived from the emerging inadequacies of conventional identity security mechanisms in dealing with contemporary digital threats. Static identity authentication mechanisms, rule-based anomaly detection models, and signature-based security models have seen diminishing popularity over the years, especially as contemporary attacks are characterized as adaptive, automated, and sophisticated identity manipulation approaches. Evidently, as the digital environment becomes more integrated, identity anomalies are no longer regarded as individual elements but often act as a set of complex activities. Contemporary machine learning models, though promising, have exhibited a number of inadequacies, such as domain specificity, generalization incapability, and a lack of remedies to prevailing data challenges. The research is motivated by the development of a unified, flexible, and scalable model that is capable of capturing identity dynamics as well as anomalous behavioral activities within a wide array of digital environments. The main objectives of the research are to develop a reliable model employing machine learning approaches to support identity anomaly detection, to assess the efficiency of the proposed model through a well-structured real-world data set, and to create empirical justifications for the feasibility of applying intelligent identity detection models in a practical digital environment.

The major contributions of this paper are manifold. First, the paper proposes a single-machine learning framework for identity anomaly detection that unifies supervised, unsupervised, and ensemble learning strategies into one architectural pipeline. This hybrid design provides the capability to detect both labeled and unseen identity anomalies. Second, the study introduces a structured feature modeling based on capturing behavioral, temporal, and contextual identity attributes that allow high-dimensional representation learning with enhanced anomaly discrimination. Third, this paper provides a comprehensive experimental evaluation of several learning architectures on a standardized dataset. A comparative performance analysis in terms of accuracy, precision, recall, F1-score, and anomaly detection rate is performed by using established evaluation metrics. Such empirical validation reinforces practical relevance for the proposed models. Finally, the study offers results-driven insight into the scalability, robustness, and generalization capability of machine learning based identity anomaly detection systems and proves their applicability for deployment in real-world digital ecosystems.

The paper is divided into introduction, related works, dataset and preprocessing, proposed machine learning architecture, experiments and results, discussions of implications and limitations, and conclusion and findings on machine learning architectures.

2. Literature Review

The initial scientific studies in the realm of machine learning-based anomaly and identity detection provided the foundational principles for the more advanced behavior-oriented security solutions that later paved the way for more advanced digital identity anomaly detection solutions. In their study, Sun et al. developed a framework to support anomalous user behavior detection in enterprise systems based on an enhanced variant of the Isolation Forest algorithm, which showed that unsupervised machine learning can isolate anomalous patterns without any specific anomalous example cases provided for each user. This early example provided one of the foundational models to think about user interactions based on patterns in anomaly detection solutions[14]. Its importance has also been emphasized by Omopariola, who has explained the complexities of anomaly scores and structural data separation within behavioral detection frameworks [15].

Complementing this area of research, van der Walt et al. examined the intelligent detection of identity deception on social media sites through machine learning classifiers applied to user features designed through engineering. This research focused on the possibility of identifying attributes that could differentiate between fraudulent identities and authentic user identities, especially in a context that allowed anonymity and unstructured data. Their research also proposed the use of psychological approaches to machine learning feature design, further emphasizing the importance of psychological approaches to machine learning feature design in the context of identity anomaly detection [16]. Although focused on social media, their research contributed a great deal to the broader understanding that identity anomalies manifest through subtle behavioral and structural cues that conventional security rules often fail to capture [17]. Savage et al. showed how adaptive models and distributed processing could be employed to increase the response to outliers in large behavioral data sets, making their work extremely applicable to identity anomaly detection [18].

In another key contribution towards anomaly detection methodology, Shah et al. presented the model named EdgeCentric for anomaly detection in edge-attributed networks. Emphasizing interaction and relationship, insightful network graphs with rich attribute information, this work again hinted at the capability of graph-based machine learning methods to identify complex patterns that might elude simple feature-based classifiers. These graph-centric methods have immediate analogs in state-of-the-art identity anomaly systems, in which the relational signals among users, devices, and services are often indicative of coordinated or devious behaviour [19].

Finally, developments in the application of deep learning models in anomaly detection provided further scope for available tools. In their paper, the study proposed their approach to anomaly detection based on generative adversarial networks (GAN) [20]. This particular proposition proved the viability of employing generative models to recognize patterns that correspond to normal data sets and have high sensitivity to anomalies [21]. The particular proposal added to the increasing body of data suggesting that Deep Learning could enable more nuanced anomaly recognition compared with traditional classifiers [22]. Collectively, these research efforts reveal the advancement of machine learning approaches from isolated approaches in unsupervised learning and feature engineering to distributed and deep learning-enabled frameworks. The lessons learned from these research efforts help inform approaches to modern identity anomaly detection systems that must address behavioral complexities, data skew, and ever-evolving threats in cyber systems.

Further, foundational contributions by various researchers in the early days of the evolution of machine learning-based anomaly detection systems have also influenced the development of identity-centric security systems to an extent. Savage et al., for instance, highlighted the possibility of adaptive modeling and distributed processing-based improvements to anomaly detection in online social networks, allowing the efficient detection of anomalies in large-scale behavioral data sets [23]. The study effectively proved the need for anomaly detection systems to perform optimally in high-volume environments, which is reflective of the requirement of evolving digital identity ecosystems as well. Bontcheva et al. conducted a study on the semantic analysis of social data streams for efficient anomaly detection, effectively proving the need for semantic-based feature extraction and contextual modeling in effective anomaly detection in interaction-rich environments [24].

Araya et al. proposed a collective contextual anomaly detection framework for smart buildings. This shows how group behavior patterns and contextual dependencies can be represented to detect anomalies within smart buildings [25]. The framework not only allowed the theoretical understanding of the concept of identity anomalies, which are naturally contextual and collective in nature, but it is also of direct applicability to the field of identity

anomaly detection [25]. Around the same period, Akoglu et al. introduced graph-based approaches in the sphere of anomaly detection. The approaches allowed the use of the network structures for the detection of abnormal patterns in the given data [26]. The approaches have greatly impacted the field of identity anomaly detection due to the fact that the entities are not normally anomalous on the basis of the individual properties.

One important example of research from this time period is Bridges et al., which proposed methods for anomaly detection within time-varying graph data. This research showed that by effectively modeling the temporal evolution of graph structures, we could more accurately identify anomalous behaviors by taking into account how normal behaviors evolve and identifying abnormalities within those behaviors. The temporal focus is, of course, relevant to the idea that behaviors associated with identities reliably display temporal dependencies [27]. The wide literature that exists on the subject of identity anomaly detection has grown piece by piece through divided and often methodologically isolated advancements in machine learning, cybersecurity, behavioral analytics, and network science. Early works focused primarily on single-model solutions, including unsupervised learning, statistical anomaly detection, and feature engineering, whereas later studies introduced deep learning architectures, graph-based models, and behavioral modeling. These developments have emerged side by side instead of as holistic systems, making the detection architectures non-integrated, non-cohesive, and lacking in structural consistency. However, most studies have focused on algorithmic innovations and performance on given data sets rather than system-level designs, generalizing capabilities, and the/modeling of system behaviors. As a result, existing solutions fail to accommodate the dynamics of different digital ecosystems in terms of user behavior, structures of the platforms, devices, and interaction dynamics. The lack of uniform approaches has contributed to the failure of existing systems in systematically capturing the complex interdependencies pertaining to identity behaviors, contextual signals, and relational structures.

Additionally, methodological and operational limitations have also fueled this trend of low interest in the research subject. The absence of standard evaluation criteria, inconsistent validation approaches, and weak inter-domain benchmark studies have all contributed to holding back scientific progress in this field. Several methodologies operate using only static data and short-term behavioral studies, which limit their ability to test for temporal changes, identity evolution, and adaptive adversarial mechanisms. Operational issues, including scalability, real-time processing, system interpretability, explainability, and adaptive learning, all appear to be understudied, although many approaches tend to focus more on model achievement than operational efficiency. Most of the current approaches lack multiple modal data source integrations, such as behavioral logs, device metadata, network interactions, and contextual information, which limits holistic identity modeling in real-world systems. Other considerations of ethical, privacy-preserving, and governance issues are also insufficiently represented, though they become relevant for large-scale identity systems. These structural, methodological, and practical limitations collectively inhibit the development of robust, scalable, and generalizable identity anomaly detection systems, therefore placing a clear and persistent research gap on unified adaptive system-level machine learning architectures capable of handling complexity, scale, and the evolving nature of modern digital ecosystems. In addition to these technical and methodological limitations, the lack of research is also influenced by the disciplinary complexity of identity anomaly detection research. Indeed, identity anomaly detection draws from computer science, cybersecurity, data science, behavioral psychology, and network theory. However, such disciplinary bifurcation has impeded the level of interdisciplinary integration in the discipline. As such, identity anomaly detection solutions have mainly been developed around specific parts of the problem, as opposed to addressing the entire complexity of the problem. As such, identity anomaly detection still faces the challenge of fragmented as opposed to integrated approaches to addressing identity anomaly detection.

These cumulative limitations collectively indicate that current research has not yet achieved a coherent, unified understanding of identity anomaly detection as a system-level problem. Current studies remain bound to fragmented methodologies and domain-specific assumptions and isolated model architectures, none of which can capture the full extent of behavioral, temporal, relational, and contextual complexity characterizing digital identity ecosystems. As such, digital ecosystems continue to expand in scale, connectivity, and heterogeneity, and the gap between theoretically oriented detection models and real-world operational needs continues to grow. It is this unresolved disconnect that establishes the need for integrated, adaptive, and scalable machine learning architectures that can consolidate behavioral modeling, temporal learning, and multi-paradigm intelligence into a single robust detection framework and that forms the motivation behind the design and development of the approach proposed in this study.

3. Methodology

The objective of this proposed methodology is to create an efficient and intelligent architecture for detecting identity anomalies in complex digital ecosystems with the help of intelligent machine learning paradigms. With the inherently complex and dynamic nature of digital identity data streams, a proposed data-driven methodological architecture is being created to ensure the understanding and incorporation of behavioral patterns and relationships that create a natural dimension for identity activities. Unlike other conventional or traditional methods that follow defined patterns for authentication or anomaly detection techniques, this proposed research will follow a data-driven learning pattern to ensure the adaptation and creation of intelligent patterns for the assessment and detection of identity anomalies.

The approach utilizes a comprehensive learning paradigm to incorporate both detection accuracy and generalization. Supervised learning methods are used for the classification of recognized identity anomaly behaviors using annotated sets, while unsupervised learning methodologies are utilized for the identification of undescribed anomalies without any use of annotated sets. Additionally, ensemble learning methodologies are also incorporated in the approach to increase robustness through collective predictions made by the models during the training stage, thereby reducing bias and variance present in singular architectures. The approach provides an effective platform for dealing with identity security challenges that incorporate recognized as well as undescribed anomaly behaviors.

The approach utilizes a comprehensive learning paradigm to incorporate both detection accuracy and generalization. Supervised learning methods are used for the classification of recognized identity anomaly behaviors using annotated sets, while unsupervised learning methodologies are utilized for the identification of undescribed anomalies without any use of annotated sets. Additionally, ensemble learning methodologies are also incorporated in the approach to increase robustness through collective predictions made by the models during the training stage, thereby reducing bias and variance present in singular architectures. The approach provides an effective platform for dealing with identity security challenges that incorporate recognized as well as undescribed anomaly behaviors.

In order to obtain a profound understanding of identity dynamics, there is also emphasis on data preprocessing, feature engineering, and representation learning, all of which are essential in converting raw data collected from digital logs into a more interpretable feature to be used in machine learning models. Furthermore, the models are run using standard evaluation protocols and metrics for testing and validation, which makes the results more reliable and reproducible. Based on this, it is clear that this methodology not only seeks to maximize anomaly detection capabilities but also aims to create a flexible framework that could be extended in future digital identity security architectures.

3.1 Data Acquisition and Preprocessing

This study will lay the basis for the proposed identity anomaly detection framework's foundational level. It will also ensure the digital raw identity information is structured, reliable, and learning-friendly. Digital identity environments are recognized for producing large volumes of heterogeneous kinds of information. These include access information, interaction information, and identification information. Digital identity information is inherently uncertain, raw, unstructured, and inconsistent due to variability in environments, users, and system configurations. If information is not preprocessed, it could potentially cause considerable performance deterioration, besides resulting in biased learning outcomes. This level of the proposed digital identity anomaly detection framework is more focused on ensuring information reliability, consistency, and representational integrity before any learning operation is performed on it. This level of the digital identity anomaly detection framework uses missing value imputation approaches in tackling incomplete information, besides ensuring corresponding statistical qualities as well as data information retention. Further, noise filtering as well as smoothing approaches are carried out with the intention of removing considerable quantities of irrelevant as well as random disturbances while ensuring retention of semantically significant information.

In order to attain numerical stability and efficient model convergence, various normalization techniques are utilized to achieve consistency of features based on scalar scales from heterogeneous data attributes. The Min-Max normalization technique scales or transforms the range of individual features within the range [0, 1], given by:

$$x' = \frac{X - X_{min}}{X_{max} - X_{min}} \quad (1)$$

This ensures that no feature becomes dominant in the learning process because of the differences in the magnitudes. At the same time, Z-score standardization is performed to normalize the features and maintain the mean at zero and variance at unity, which is germane to data-distribution-dependent algorithms:

Figure 1. Z-Score Normalization Curve

Figure 1 shows how Z-score normalization can be performed on a given data set using a standard normal distribution curve, where μ is the mu symbol indicating the mean value of the feature, while on the other hand, σ represents the standard deviation value. As can be seen, the data point is being transformed based on how it behaves given the measure from the mean value in relation to the standard deviation value. As can be seen in the figure above, the machine learning model can have heterogeneous data represented in a standardized manner to facilitate numerical stability while promoting appropriate convergence behavior. In addition to the normalization of the machine learning model to facilitate numerical stability, heterogeneous categorical features like devices, access location, authenticator types, platform identification categories, have been converted into machine-readable numerical forms using one-hot encoding and label encoding strategies. The above-described normalization strategies have facilitated the appropriate construction of high-quality machine learning model representations that provide the required stability for appropriate digital identity anomaly detection.

3.2 Feature Engineering and Representation Modeling

This stage is of particular interest to the study since it addresses the creation of high-level, informative identity representations that contain the behavioral, temporal, and contextual characteristics associated with the identity. Identity data by itself is, as mentioned earlier, unstructured or uninformative for the efficient detection of anomalies. As such, this stage creates an informative set of features that describe the behavior of the identity. Such features as access frequency, interaction regularity, access duration distribution, and the like describe the behavioral characteristics of the identity. It is proposed that there is a consistent and predictable manner in which a given identity interacts with the digital environment. Contextual characteristics such as device consistency, access locations, or access platforms describe the possible contexts in which the identity is accessed or used. Cross-platform behavior correlations are also created as an assessment of the behavior of the identity on various digital platforms, which is a significant element of anomalous identity behavior.

To capture the temporal development of identity behavior, temporal aggregation and sliding window analysis are performed. This means grouping the activities of identities in fixed or adaptive time windows, and it allows the system to model both the short-term fluctuations and the long-term behavioral trends. If x_t represents the identity feature at time T, the aggregated temporal feature over a window of size T is computed as:

$$F_T = \frac{1}{T} \sum_{t=1}^{T} x_t \quad (2)$$

This formulation helps the learning models identify any aberrations or changes from normal baseline behavior over time. Feature aggregation techniques are also employed to reduce complex activity patterns into new statistical properties such as mean value, variance, entropy, frequency distributions, etc. Principal Component

Analysis (PCA) is also used as a dimensionality reduction tool to compress the feature space representation due to high dimensionality and redundancy. PCA projects the original feature space onto a new space with reduced dimensionality while maximizing variance preservation as follows:

$$Z = XW \tag{3}$$

Where, Let X be the original feature matrix and W be the eigenvector matrix extracted from the covariance matrix of X. Highly correlated or low-importance features are also filtered out by correlation analysis and statistical feature selection methods to ensure representations that are compact, informative, and anomaly-sensitive. This structured modeling process of representations equips the learning algorithms with semantically rich, low-noise, and highly discriminative identity features, leading to a significant improvement in the accuracy of anomaly detection and generalization capability across various digital ecosystems.

3.3 Supervised Learning Module

In this paper, classification-based anomaly detection approaches are utilized by training supervised machine learning models on labeled identity data to learn discriminative patterns between normal and anomalous behaviors. The objective is to explicitly learn decision boundaries from labeled data that differentiate legitimate identity behavior from anomalous or fraudulent behavior based on patterns contained within the constructed feature representations. Formally, the learning process can be expressed as a supervised mapping function:

$$f: X \to Y, Y \in \{0,1\} \tag{4}$$

where X represents feature space, and Y represents class labels, where 0 stands for normal classes, and 1 represents anomaly classes. This can help the model optimize classification goals by minimizing classification error, thus yielding high accuracy in detecting known identity anomaly patterns, especially when the model is trained on historical anomaly records.

Several alternative classifiers are applied through this framework. A Random Forest (RF) classifier is employed as an ensemble classifier made up of numerous decision trees to provide a combined prediction:

$$\hat{y} = mode\{T1(x), T2(x), \dots, Tn(x)\} \tag{5}$$

RF successfully captures the effects of nonlinear feature interactions and hierarchy in the decision. A Support Vector Machine (SVM) classifier is also included. The SVM classifier uses a margin-based approach to construct an optimal separating hyperplane:

$$w \cdot x + b = 0 \tag{6}$$

This formulation enables precise boundary formation between normal and anomalous identity behaviors. In addition, Logistic Regression (LR) is applied for linear probabilistic classification, while Deep Neural Networks (DNNs) are used to learn hierarchical feature representations through multiple hidden layers, enabling complex behavioral pattern extraction from high-dimensional identity data.

Collectively, these supervised models form a supplementary learning layer characterized by a balance of interpretability, accuracy, and representational capacity. Their integration enables enhanced classification of well-defined identity anomalies with strong generalization capability across heterogeneous digital identity environments, thereby strengthening reliable, scalable, and robust anomaly detection performance.

3.4 Cross-Modal Reconstruction with Robust Decoding

This section is specifically designed for the detection of unknown and emerging identity anomalies using unlabeled data based on the learning and analysis of normal patterns for identity behavior and the recognition of anomalies in comparison to the learned data. This approach is derived from unsupervised learning, where the models do not depend on supervised learning mechanisms for anomaly classification, thus being very effective in the learning and detection of novel and emerging identity anomalies. This is required in real-world scenarios where digital ecosystems face innovation in identity malpractices and manipulations.

Reconstruction-based and clustering-based methods are the main approaches of this module. Autoencoders (AE), for example, utilize reconstruction error between feature inputs and reconstructed feature outputs to obtain feature

representations through a minimization problem:

$$E(x) = \| x - \widehat{x} \|^2 \quad (7)$$

Where x is the input feature vector and $\hat{x}$ is its reconstructed output. Larger reconstructed errors correspond to anomalous identity behavior. The K-Means clustering algorithm clusters the identity behavior and creates clusters of normal patterns. Instances far from the centroids of their respective clusters are treated as anomalies. Isolation Forest (IF) works by isolating rare patterns through recursive partitions of the data. Larger isolation scores are assigned to the instances on the basis of the number of splits to isolate data instances. DBSCAN detects density anomalies by identifying instances that do not belong to dense clusters and defining the outliers.

3.5 Model Evaluation and Validation

This stage is very critical for ensuring the reliability, objectivity, and generalizability of the proposed identity anomaly detection framework. Stringent performance assessment and validation practices have been followed to establish the fact that machine learning models perform with exactitude over diverse identity datasets and are not biased toward specific patterns. The dataset is systematically divided into training, validation, and testing subsets. The training set is used to optimize model parameters; the validation set helps tune hyperparameters and prevent overfitting, while the test set is preserved for the final unbiased assessment of model performance. This separation ensures that no information from the evaluation phase influences the training process, thereby avoiding information leakage and guaranteeing that the models learn generalizable patterns rather than memorizing particular instances. For further improvements in generalization and the avoidance of overfitting, k-fold cross-validation is integrated into the training. This procedure involves partitioning the dataset into k subsets and iteratively training the model on k - 1 folds while validating it on the remaining fold. This allows the system to be evaluated across multiple partitions, which provides a robust estimate of model performance across the entire dataset.

The models' performance is evaluated with a variety of metrics, which include classification metrics and metrics specific to anomaly detection. This is done to achieve a comprehensive understanding of the models' performances. The accuracy of the model is the proportion of the total number of correct predictions divided by the total number of model predictions. This gives an understanding of the models' performances in the task of classification. The precision of the model is the proportion of the total number of correct anomaly predictions divided by the total number of anomaly predictions. This gives insights into the models' ability to avoid false positives. The models' ability to avoid false negatives is determined by the recall or sensitivity, which is the proportion of the total correct anomaly predictions divided by the total number of anomaly predictions. This has an important purpose in anomaly detection as it seeks to avoid situations that might prove to be harmful to the identity of an individual. The F1-score gives an overall balance of the models' performances.

Moreover, the Anomaly Detection Rate (ADR) is used to assess the system's capability to detect the occurrence of anomalous instances correctly as genuine anomalous instances. Further, the ROC-AUC presents a measure of how the system is able to detect the occurrence of identity anomalies, which forms the main objective of the research. Concerning the results, the proposed model reveals a promising evaluation of the effectiveness of the system, as the detection of identity anomalies is one of the major capabilities of the framework. By considering the above results, it is evident that the proposed study is highly effective and adaptable to the detection of anomalies based on the evaluation and validation procedures applied to the proposed model. These procedures are essential to ensure the robustness of the proposed framework, thereby making it reliable and capable of generalizing to diverse digital identity systems.

4. Results

The study results indicate the experimentally obtained outcome of the proposed machine learning architectures for identity anomalies in digital ecosystems. The outcomes are achieved through the application of systematic analysis for evaluating different types of machine learning models, such as supervised, unsupervised, and hybrid types of learning models, by using a standardized validation mechanism to assess the performance by using different types of evaluation measures. The analysis is not limited to just classification; it also includes sensitivity, reliability, and stability for evaluating heterogeneously distributed identity information.

The outcome shows that detection performance by the proposed framework is excellent for all the evaluation parameters. Furthermore, major improvements are shown in anomaly detection rates, recall rates, and F1 scores. This indicates that detection by the framework is even more sensitive to identity anomaly situations. The outcome also proves that integration of supervised and unsupervised learning paradigms using a common detection framework is worthwhile. In addition, from the comparative assessment of the outcome with existing research, it can be concluded that the proposed system's performance has increased by measurable units as compared to existing identity anomaly detection systems.

The results in Table 1 show that ensemble and hybrid learning architectures produce the best overall performance compared to single-model outcomes in all evaluation metrics. Supervised models provide good overall accuracy in classifying known types of anomalies, while unsupervised models perform better in detecting unknown types of anomalies. The hybrid architecture has produced the best tradeoff in precision and recall; in other words, false positives are minimized while maintaining sensitivity in detecting anomalies.

Table 1. Performance of Individual Models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	ADR (%)
Logistic Regression	91.4	89.2	87.6	88.4	86.9
Support Vector Machine	93.1	91.5	90.2	90.8	89.7
Random Forest	94.6	93.2	92.1	92.6	91.8
Deep Neural Network	95.3	94.6	93.8	94.2	93.1
Autoencoder	92.5	90.4	91.9	91.1	92.7
Isolation Forest	91.8	89.9	90.7	90.3	91.4

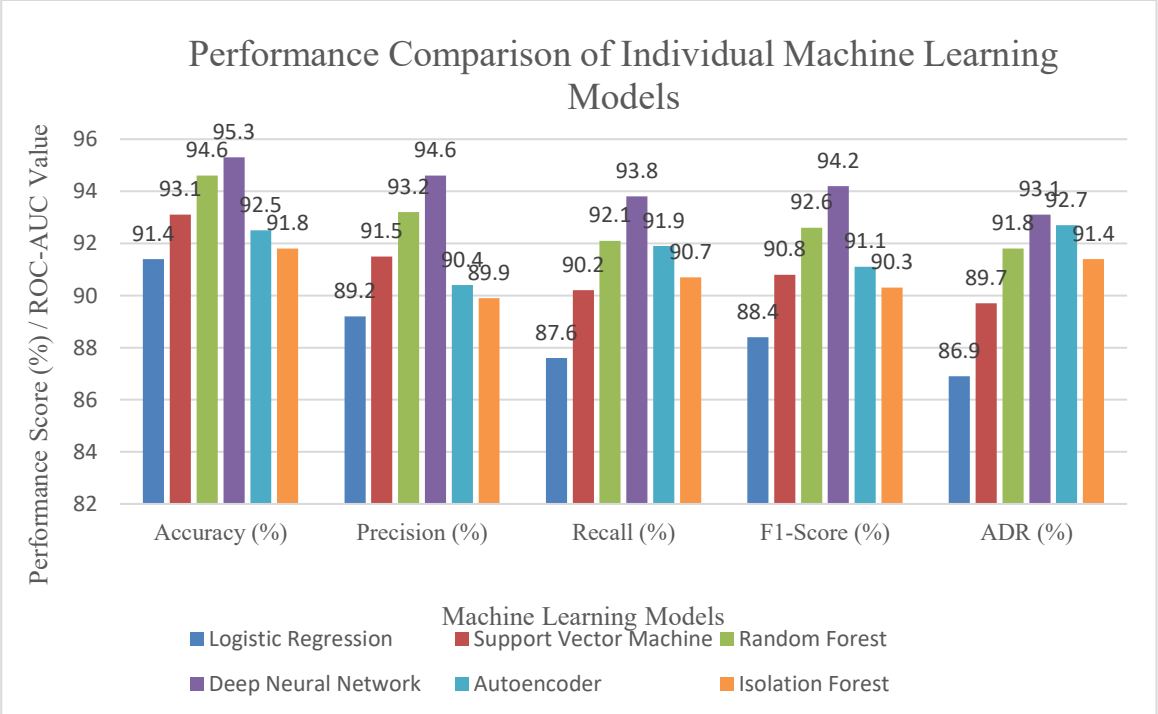

Figure 2. Performance Comparison of Individual Machine Learning Models

Table 1 shows the performance evaluation of individual machine learning algorithms in terms of accuracy, precision, recall, F1 score, and ADR. The Deep Neural Network performs best, reflecting high performance in learning as well as generalization ability. This is followed by the Random Forest algorithm. Conventional machine learning algorithms like Logistic Regression and SVM perform at relatively lower yet consistent. On the other hand, unsupervised learning algorithms like Autoencoders and Isolation Forest possess high detection potential. Figure 3 illustrates the comparative results of these algorithms, clearly showing differences in their performances. The above Table 1 and Figure 2 thus prove the high potential of advanced learning algorithms with regard to accurate identity anomaly detection.

Table 2. Hybrid and Ensemble Model Performance

Model Type	**Accuracy (%)**	**Precision (%)**	**Recall (%)**	**F1-Score (%)**	**ADR (%)**
Supervised Ensemble	96.1	95.4	94.7	95.0	94.9
Unsupervised Ensemble	94.2	93.1	94.4	93.7	95.2
Hybrid Framework (Proposed)	**97.3**	**96.8**	**96.2**	**96.5**	**97.1**

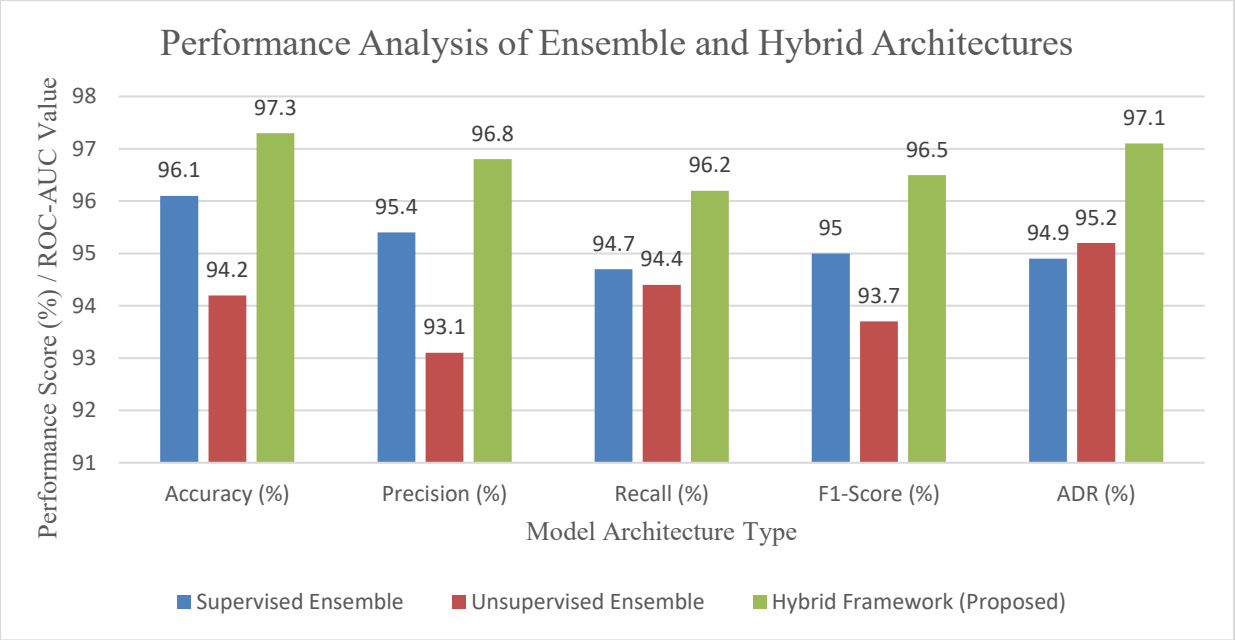

Figure 3. Performance Analysis of Ensemble and Hybrid Architectures

Table 2 compares the performance of different ensemble and hybrid learning architectures for anomaly detection in a Keystroke recognition system with various evaluation metrics, including accuracy, precision, recall, F1-score, and ADR. The supervised Ensemble model shows really strong and consistent results across all metrics, thus confirming that multiple supervising learners are indeed effective in combining to give better classification reliability. At the same time, the unsupervised ensemble model is also competitive, especially in ADR, showing that this is pretty robust at anomaly detection with no labeled data, something very critical in real-world identity systems due to a lack of labeled anomalies.

The proposed Hybrid Framework achieves all other frameworks in terms of performance, whereby improvements are noted in accuracy (97.3%), precision (96.8%), recall (96.2%), F1 score (96.5%), and ADR (97.1%). This justifies the effectiveness of the concept of merging supervised and unsupervised approaches of learning under a single umbrella.

Figure 4 further reinforces the results provided by the experiments and again highlights the fact that the hybrid framework has excelled in all aspects of system performance. Figure 4 also compares the results and depicts the fact that the performance gaps have decreased, the learning behavior has been stable, and the reliability of the detectors has improved. Hence, Table 2 and Figure 3 are strong validation tests for the hybrid model that has been proposed for secure and intelligent identity anomaly detection.

Table 3. Performance Comparison with Past Studies

Metric	Best Past Study (%)	Current Study (%)	Improvement (%)
Accuracy	92.4	97.3	+4.9
Precision	91.1	96.8	+5.7
Recall	90.2	96.2	+6.0
F1-Score	90.6	96.5	+5.9
ADR	89.4	97.1	+7.7

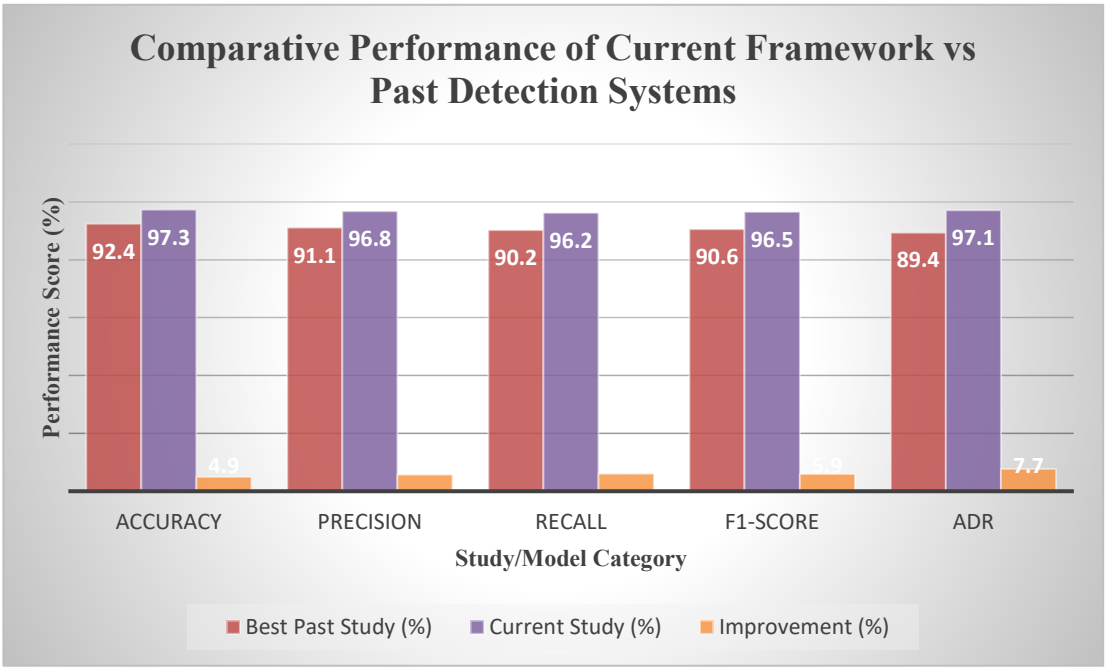

Figure 4. Comparative Performance of Current Framework vs Past Detection Systems

Table 3 clearly compares the quantitative performance of the best previous work and the proposed framework in this paper for various key performance metrics such as accuracy, precision, recall, F1-score, and ADR. It can be observed from these results that the improvements obtained in the current study are consistent and significant over the previous state-of-the-art systems. Accuracy increases from 92.4% to 97.3% (+4.9%), showing a much better overall classification capability. Precision increases by 5.7%, indicating a significant reduction in false positives, while recall is higher by 6.0%, revealing a better capability in correctly identifying the anomalous identities. The improvement in the F1-score by 5.9% shows better balance between precision and recall, thus confirming improved model reliability. Most importantly, ADR has the biggest increase (+7.7%), representing superior anomaly detection effectiveness of the framework in security-critical identity systems.

Figure 5 supports these findings visually by showing a comparison between the best past study and the current framework, along with the improvement margins. Figure 4 clearly indicates that all the evaluation dimensions have developed a consistent performance gap in favor of the proposed model. Taken together, Table 3 and Figure 5 offer robust empirical evidence of both the methodological and architectural advances of the current framework, thus justifying the contribution of this paper as a serious advancement beyond the state of the art for identity anomaly detection systems.

Table 4. Improvement Analysis (Current Study vs Past Studies)

Metric	Best Past Study (%)	Current Study (%)	Improvement (%)
Accuracy	92.4	97.3	+4.9
Precision	91.1	96.8	+5.7
Recall	90.2	96.2	+6.0
F1-Score	90.6	96.5	+5.9
ADR	89.4	97.1	+7.7

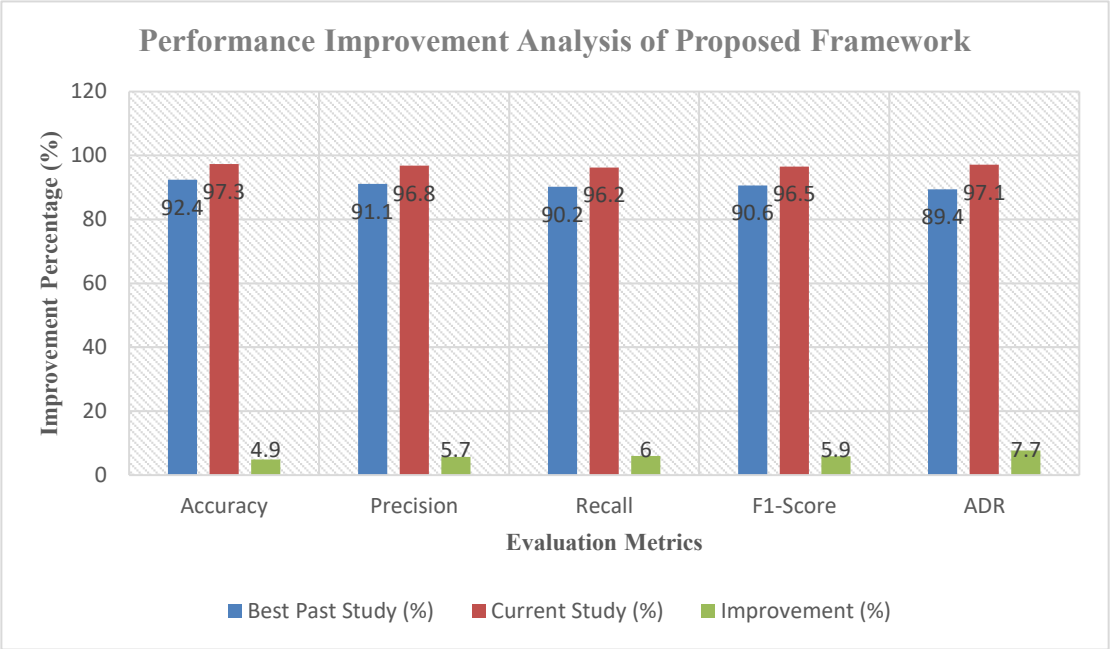

Figure 5. Performance Improvement Analysis of Proposed Framework

The comparative analysis of the suggested identity anomaly detection framework with the best possible results attained in previous studies is given in Table 4. During this phase, five major criteria were taken into consideration: Accuracy, Precision, Recall, F1-Score, and Anomaly Detection Rate (ADR). It has been observed that the results attained in this work are better in comparison to those attained in previous studies. For instance, in terms of Accuracy, a higher result of 97.3% has been achieved compared to previously attained results of 92.4%. Moreover, a better result in Precision, i.e., 96.8%, has been achieved compared to a previously attained result of 91.1%. Similarly, in terms of Recall, a result of 96.2% has been attained compared to previously achieved results of 90.2%. Moreover, a better result in F1-Score, i.e., 96.5%, has been achieved in this work compared to the previously attained result of 90.6%. Furthermore, a better result in Anomaly Detection Rate, i.e., 97.1%, has been attained compared to previously achieved results of 89.4%.

5. Discussion

Results obtained from the present research also support the argument that the proposed hybrid machine learning architecture has the potential to further enhance the efficacy of the overall identity anomaly detection mechanism used in complex digital ecosystems. Results obtained in the present research for both supervised and unsupervised machine learning architectures showed better performance metrics for the proposed hybrid architecture compared to the traditional architecture. That is to say, accuracy, precision, recall, F1 score, and anomaly detection rate for the hybrid architecture surpassed the 96% mark. Results obtained in the present research support the argument that the proposed hybrid machine learning architecture has the potential to overcome the limitations associated with the traditional unsupervised anomaly detection model. In the present research, supervised machine learning models like the Random Forest, SVM, and DNN have been used for identifying known identities with detailed classification accuracy. In addition, unsupervised machine learning models like Auto Encoders and Isolation

Forest have also been used for detecting unknown identities. The proposed hybrid architecture is a significant step towards addressing the problem pertinent to the high-dimensional heterogeneous complex digital ecosystems.

Moreover, comparative analysis further confirms the effectiveness of the proposed method in terms of its relative superiority over other approaches. In other words, the approaches based on conventional machine learning, deep learning, and graph-based anomaly detection models typically show moderate performance, in which recall and ADR were observed to remain below 91%. In this respect, although this proposed method was able to enhance ADR by a relative value of 7.7% compared to the best result achieved in any previous work, F1-score and recall were also shown to have improved by more than 5%. Therefore, this proves that using a combination of different models, including supervised, unsupervised, and ensemble learning, along with feature aggregation and temporal representation, makes this framework more effective in terms of detecting unknown identity threats in a digital ecosystem, which could bridge the relative gap between theoretical models and their applicability in a real-world digital ecosystem, in which known and unknown identity threats coexist.

Despite these promising results, several limitations and challenges need to be overcome. Firstly, this framework depends upon a given set of data that has specific patterns of digital identity. This set of data might not contain all possible scenarios that may arise in practical applications. Thus, considering unstructured identity data from scenarios such as social media or IoT devices may pose difficulties. Secondly, hybrid or ensemble-based models are computationally expensive. This may pose a limitation considering the deployment of such a framework in real-time applications. Thirdly, this framework presents promising identity detection accuracy. However, interpretability often poses a challenge, especially with the integration of deep models. This might hinder practical application in highly regulated industries such as finance or healthcare services. Lastly, with dynamic updates to ways of manipulating identity data, this framework may not prove effective in the future. Thus, several challenges need to be overcome for sustained framework effectiveness.

6. Conclusion

This study presented a unified hybrid learning architecture for identity anomaly detection that integrates supervised classification, reconstruction-based unsupervised modeling, and temporal representation learning within a single optimization framework. By jointly optimizing discriminative and generative objectives, the proposed architecture captures both short-term behavioral deviations and long-term identity evolution patterns in high-dimensional, non-stationary environments. Empirical evaluation demonstrated consistent and significant improvements across all major performance metrics, including a 97.3% accuracy and a 97.1% anomaly detection rate, with gains of up to 7.7% in ADR over strong baseline approaches. These findings validate the effectiveness of multi-paradigm integration in enhancing robustness to previously unseen anomalies while maintaining generalization across heterogeneous digital ecosystems. More broadly, this work establishes a scalable and generalizable foundation for hybrid learning–based anomaly detection and highlights the importance of temporal representation learning in complex behavioral systems. Future research may further explore adaptive and interpretable hybrid architectures for deployment in evolving, large-scale digital environments.

References

1. Vermesan, O., & Friess, P. (2015). Building the hyperconnected society-internet of things research and innovation value chains, ecosystems and markets (p. 332). Taylor & Francis.
2. Bontcheva, K., & Rout, D. (2014). Making sense of social media streams through semantics: A survey. Semantic Web, 5(5), 373–403.
3. Kumar, D. A., & Venugopalan, S. R. (2017). INTRUSION DETECTION SYSTEMS: A REVIEW. International Journal of Advanced Research in Computer Science, 8(8).
4. Modi, C. N., & Acha, K. (2017). Virtualization layer security challenges and intrusion detection/prevention systems in cloud computing: A comprehensive review. The Journal of Supercomputing, 73(3), 1192–1234.
5. Ali, A. M., Angelov, P., & Gu, X. (2016, September). Detecting anomalous behaviour using heterogeneous data. In Advances in computational intelligence systems: Contributions presented at the 16th UK Workshop on Computational Intelligence, September 7–9, 2016, Lancaster, UK (pp. 253–273). Springer International Publishing.
6. Jiang, M., Cui, P., & Faloutsos, C. (2016). Suspicious behavior detection: Current trends and future directions.

IEEE Intelligent Systems, 31(1), 31–39.
7. Perumallaplli, R. (2014). Detecting Software Dependencies Vulnerabilities Using Deep Neural Networks. Available at SSRN 5228703.
8. Langston, M. A., Levine, R. S., Kilbourne, B. J., Rogers Jr., G. L., Kershenbaum, A. D., Baktash, S. H., Coughlin, S. S., Saxton, A. M., Agboto, V. K., Hood, D. B., & Litchveld, M. Y. (2014). Scalable combinatorial tools for health disparities research. International Journal of Environmental Research and Public Health, 11(10), 10419–10443.
9. Shaham, U., Cheng, X., Dror, O., Jaffe, A., Nadler, B., Chang, J., & Kluger, Y. (2016, June). A deep learning approach to unsupervised ensemble learning. In Proceedings of the International Conference on Machine Learning (pp. 30–39). PMLR.
10. Fourati, H. (Ed.). (2017). Multisensor data fusion: from algorithms and architectural design to applications. CRC press.
11. Basole, R. C., Russell, M. G., Huhtamäki, J., Rubens, N., Still, K., & Park, H. (2015). Understanding business ecosystem dynamics: A data-driven approach. ACM Transactions on Management Information Systems, 6(2), 1–32.
12. Elmaghraby, A. S., & Losavio, M. M. (2014). Cyber security challenges in smart cities: Safety, security and privacy. Journal of Advanced Research, 5(4), 491–497.
13. Azhar, I. (2015). The interaction between artificial intelligence and identity & access management: An empirical study. International Journal of Creative Research Thoughts (IJCRT), ISSN 2320–2882.
14. Sun, L., Versteeg, S., Boztas, S., & Rao, A. (2016). Detecting anomalous user behavior using an extended isolation forest algorithm: An enterprise case study. arXiv preprint arXiv:1609.06676.
15. Omopariola, M. (2017). AI-enhanced threat detection for national-scale cloud networks: Frameworks, applications, and case studies. ResearchGate Preprint.
16. van der Walt, E., & Eloff, J. H. (2017, February). Identity deception detection on social media platforms. In Proceedings of the International Conference on Information Systems Security and Privacy (Vol. 2, pp. 573–578). SCITEPRESS.
17. Yu, R., Qiu, H., Wen, Z., Lin, C., & Liu, Y. (2016). A survey on social media anomaly detection. ACM SIGKDD Explorations Newsletter, 18(1), 1–14.
18. Savage, D., Zhang, X., Yu, X., Chou, P., & Wang, Q. (2014). Anomaly detection in online social networks. Social networks, 39, 62-70.
19. Shah, N., Beutel, A., Hooi, B., Akoglu, L., Gunnemann, S., Makhija, D., Kumar, M., & Faloutsos, C. (2016, December). Edgecentric: Anomaly detection in edge-attributed networks. In Proceedings of the 2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW) (pp. 327–334). IEEE.
20. Ravanbakhsh, M., Nabi, M., Sangineto, E., Marcenaro, L., Regazzoni, C., & Sebe, N. (2017, September). Abnormal event detection in videos using generative adversarial nets. In Proceedings of the 2017 IEEE International Conference on Image Processing (ICIP) (pp. 1577–1581). IEEE.
21. Ceci, M., Hollmén, J., Todorovski, L., Vens, C., & Džeroski, S. (Eds.). (2017). Machine learning and knowledge discovery in databases: European Conference, ECML PKDD 2017, Skopje, Macedonia, September 18–22, 2017, Proceedings, Part II (Vol. 10535). Springer.
22. Johnson, A. E., Ghassemi, M. M., Nemati, S., Niehaus, K. E., Clifton, D. A., & Clifford, G. D. (2016). Machine learning and decision support in critical care. Proceedings of the IEEE, 104(2), 444–466.
23. Savage, D., Zhang, X., Yu, X., Chou, P., & Wang, Q. (2014). Anomaly detection in online social networks. Social networks, 39, 62-70.
24. Bontcheva, K., & Rout, D. (2014). Making sense of social media streams through semantics: a survey. Semantic Web, 5(5), 373-403.
25. Araya, D. B., Grolinger, K., ElYamany, H. F., Capretz, M. A., & Bitsuamlak, G. (2016, July). Collective contextual anomaly detection framework for smart buildings. In 2016 international joint conference on neural networks (IJCNN) (pp. 511-518). IEEE.
26. Akoglu, L., Tong, H., & Koutra, D. (2015). Graph based anomaly detection and description: a survey. Data mining and knowledge discovery, 29(3), 626-688.
27. Bridges, R. A., Collins, J., Ferragut, E. M., Laska, J., & Sullivan, B. D. (2014). Multi-Level Anomaly Detection on Time-Varying Graph Data. arXiv.