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

DeepSynth: A Robust Multi-Layer Neural Detection of Coordinated Latent Anomalies in High-Dimensional Identity Systems

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

Submitted: 17/12/2018 **Revised**: 20/01/2019 **Accepted**: 16/02/2019

Abstract: The proliferation of high-dimensional, multi-modal behavioral signals in modern digital identity ecosystems has been accompanied by the evolution of adversarial strategies from isolated point anomalies to coordinated deviations across correlated feature subspaces. Existing anomaly detection frameworks primarily model marginal feature deviations, which limits their sensitivity to the higher-order dependency structures that characterize coordinated attacks. This research introduces DeepSynth, a coordination-aware multi-layer neural framework to detect synchronized latent anomalies in complex identity ecosystems. The framework integrates hierarchical representation learning with covariance-informed latent modeling to explicitly capture inter-feature dependencies. Anomaly scoring integrates reconstruction residuals, Euclidean latent deviation, Mahalanobis distance, and a normalized covariance-based coordination metric to quantify coordinated deviations across latent features. A latent-level ensemble aggregation coupled with theoretical variance-reduction enhances robustness against noise, imbalance, and adversarial variability. Empirical evaluation on heterogeneous identity datasets comprising behavioral logs, transactional records, and synthetic attack simulations demonstrates the efficacy of the framework. DeepSynth achieves a peak detection accuracy of 93.8% and an AUC-ROC of 95.1%, significantly outperforming strong baselines including Deep SVDD, Isolation Forest, and LSTM encoder-decoders ($p < 0.01$, paired bootstrap test). Furthermore, component-wise analysis confirms that explicit latent coordination modeling provides measurable gains beyond reconstruction-only objectives. These findings establish that resolving higher-order latent dependency structures is critical for robust anomaly detection in high-dimensional identity ecosystems. Future work includes deriving theoretical generalization bounds for coordination-aware anomaly detection under evolving covariance structures.

Keywords: *Coordinated anomaly detection, Deep neural networks, Ensemble robustness, High-dimensional data, Identity systems, Latent representation learning*

1. Introduction

The emergence of digital identity platforms has created a set of complex issues in user management, authentication, and tracking on the internet. The platforms are now using high-dimensional data such as behavioral patterns, biometric data, and transactional patterns, making it difficult to identify irregular or malicious activities [1]. Conventional machine learning algorithms have been used to identify anomalies, mainly concentrating on individual anomalies in a single feature or a combination of features [2]. These approaches can identify anomalies but are not capable of identifying coordinated anomalies that occur in multiple dimensions simultaneously [3]. Coordinated anomalies can occur in fraud, identity attacks, or collusions among multiple accounts, which are serious security threats and cannot be identified using shallow models. The complexity and high dimensionality of identity data require sophisticated models that can identify hidden patterns and relationships, which can provide effective anomaly detection without depending on simple rules or heuristics [4].

Recent breakthroughs in representation learning and deep neural networks have shown the potential to uncover hidden features from high-dimensional data, facilitating more precise modeling of normal and anomalous behavior. Research has emphasized the importance of hierarchical models for learning multi-layered representations, which can uncover subtle patterns that are not captured by conventional methods [5]. Moreover, research has investigated the benefits of ensemble and hybrid models, which integrate different learning paradigms to improve robustness against noise and adversarial attacks. In the context of digital identity management systems, the capability to uncover coordinated and hidden anomalies has been identified as a key need, as attackers have been increasingly using sophisticated methods that exploit correlations across multiple accounts or dimensions of behavior [6]. The confluence of high-dimensional feature spaces, complex behavioral dependencies, and adversarial methods has thus driven the need for deep, multi-layered models that can uncover hidden anomalies while preserving generalization and noise robustness [7]. Overall, this research literature provides the basis for developing methods that combine deep learning, feature

[1*,2]*Department of Computer Science*
California State University, East Bay,
Hayward, USA
** Corresponding Author Email:*
ssanjeevprasanna@horizon.csueastbay.edu

representation, and anomaly detection to tackle the challenges of secure, large-scale identity management systems [8].

The proposed research aims at developing a strong deep learning solution for coordinated latent anomaly detection in high-dimensional identity systems. The solution is intended to identify latent anomalies that are not easily detected by traditional machine learning techniques, especially when the anomalies are scattered over various features or accounts. This research is particularly important in today's digital identity systems, where behavior, biometric, and transactional data co-exist in complex high-dimensional spaces that require hierarchical representation learning. The proposed research is driven by the increasing use of digital identity systems and the sophistication of coordinated attacks. The research emphasizes accuracy, robustness, and scalability. The main contributions of the proposed research are the development of a coordinated latent anomaly detection model, the development of a covariance-based coordination measure in the latent space, the development of an ensemble-driven robustness module for dealing with noise and adversarial behavior, and experimental validation. The proposed research is divided into background, methodology, experimental evaluation, and analysis.

2. Literature Review

The literature on anomaly detection is characterized by a wide range of approaches to detect irregular patterns in data that do not conform to the learned norm. There is a substantial amount of literature on deep learning-based approaches that have the ability to capture complex patterns in high-dimensional data that are difficult to model using conventional approaches. The use of hierarchical models and generative models is effective in detecting anomalies in challenging conditions, such as semi-supervised learning and high-dimensional data. The need for interpretable and efficient models has also been highlighted, resulting in improvements in model architecture and evaluation metrics [9].

The research work of N. Souly et al. [10] introduces a semi-supervised method that employs a conditional generative adversarial network to achieve the simultaneous learning of high-dimensional data generation and its latent representation, which facilitates better anomaly detection with one-class training paradigms; the encoder-decoder-encoder structure of the model enables the anomalies to be discriminated based on large differences in the learned latent distributions and has been proven on various benchmark datasets. The paper by Yasuhiro Ikeda et al. [11] proposes a new algorithm based on VAE that aims to detect the dimensions with the most influence on anomaly decisions in high-dimensional data, thus enhancing the interpretability of the latent space in anomaly detection problems by estimating the dimensions of the features that contribute to the detection.

In the study by Houssam Zenati et al. [12] An adversarially learned anomaly detection system framework using bi-directional GANs is proposed, which derives adversarially learned features for anomaly detection and leverages reconstruction errors in both data space and latent space to improve performance and stability in high-dimensional anomaly detection tasks. The study by Lim et al. [13] (DOPING: Generative Data Augmentation for Unsupervised Anomaly Detection with GAN) applies generative models to oversample the normal instances in multimodal high-dimensional distributions to counter false positives and improve the robustness of the detection process through adversarial autoencoder techniques to map distributions to latent spaces for better unsupervised anomaly detection.

In the research conducted by Y. Thu et al. [14], improvements in the adversarial autoencoder methods are highlighted, specifically on the encoder-decoder-encoder architecture designed for anomaly problems, emphasizing the importance of learning the latent space and adversarial training in distinguishing normal and anomalous patterns in complex data environments. Although it is not a deep learning-centric paper, the paper by J. Ker et al. [15] reviews deep learning techniques for anomaly detection and classifies them into feature extraction, normality representation, and end-to-end score learning models, emphasizing the superiority of deep learning models over shallow models in modeling complex relationships for anomaly detection models.

Ruff et al. [16] proposed a deep one-class classification method that reformulated the anomaly detection task as a hypersphere learning problem in the latent space. The proposed method, known as Deep SVDD, learns to represent data points in a compact hypersphere that defines normal data, and any point outside this hypersphere is considered an anomaly. The proposed method is a significant shift from shallow one-class classification to deep representation-based anomaly detection. The proposed method is robust to noise and irrelevant feature variations due to the enforced compact latent representation. However, the proposed method assumes that anomalies are point-wise deviations in the global latent space, which may not always be the case when anomalous behaviors are distributed in multiple dimensions. This assumption becomes more apparent in identity systems, where anomalies may look normal individually but anomalous collectively.

Zong et al. [17] introduced the Deep Autoencoding Gaussian Mixture Model, which combines deep autoencoders with probabilistic density estimation. This

approach enables unsupervised anomaly detection by leveraging the strengths of deep autoencoders and probabilistic density estimation. The model learns low-dimensional representations and models their distribution using Gaussian mixture components. This approach is more effective than traditional reconstruction-based models because it leverages distributional information. However, the model has limitations because it is mainly focused on modeling the marginal distribution of the latent variables and does not model the dependency relationships among the variables. This means that coordinated anomalies based on multiple identity variables might not be captured accurately.

Zhou and Paffenroth [18] proposed robust deep autoencoders to overcome the weakness of traditional autoencoders against corrupted data and outliers. Their method integrates robustness constraints into the reconstruction process, which helps to distinguish the structure of normal data from sparse anomalies. This method showed robust performance in high-dimensional and noisy settings, and it applies to real-world anomaly detection tasks. Nevertheless, the robustness technique is mainly designed for sparse and independent corruptions, but not for structured or coordinated anomalies. In identity-driven systems, anomalous behaviors usually change in a correlated manner across features or entities, which is difficult for reconstruction models focused on robustness to capture.

Marchi et al. [19] offered one of the early deep learning solutions for anomaly detection via denoising autoencoders and temporal modeling of acoustic signals. The solution learns to represent normal temporal patterns and detects anomalies based on the reconstruction error over time. This solution showed the promise of deep learning models to encode dependencies in sequential data and temporal patterns, going beyond the boundaries of static feature analysis. Although successful in temporal novelty detection, the solution is highly dependent on the reconstruction error and lacks the ability to capture higher-order relationships in the latent spaces. Consequently, its applicability to complex identity systems with coordinated, multi-attribute anomalies remains limited.

Kim et al. [20] used long short-term memory networks for multivariate anomaly detection in intrusion detection problems. The method is useful for modeling the temporal relationships between multiple features, thus improving the detection of evolving anomalous behavior that evolves. This paper demonstrated the significance of temporal relationships in anomaly detection and the benefits of using recurrent neural networks over static models. Nevertheless, the emphasis is still

on sequence-level anomalies and not on the coordination of latent structures in feature subspaces. The model does not make use of or rely on the coordinated latent structures, which are essential for identity anomaly detection that occurs over multiple accounts or dimensions.

Table 1. Summary of Key Anomaly Detection Studies

Study	Methods	Key Findings
[21]	Hybrid DBN + One-Class SVM deep model for high-dim anomaly detection	Combines deep belief networks with one-class SVM to extract robust features and detect anomalies in high-dimensional data, improving efficiency and accuracy vs. standalone methods.
[22]	LSTM-based encoder-decoder for multi-sensor anomaly detection	Uses LSTM encoder-decoder to reconstruct normal behavior and detect anomalies via reconstruction error in diverse time-series datasets, proving robustness to unpredictable patterns. Paper link:
[23]	LSTM + OC-SVM / SVDD for sequence anomaly detection	Proposes joint training of LSTM with one-class SVM or SVDD to detect anomalies in variable-length sequence data, enhancing performance in unsupervised contexts.
[24]	DeepNet for appearance & motion anomaly detection	Introduces deep unsupervised learning to automatically extract features from both appearance and motion data for event anomaly detection in complex video scenes
[25]	Contractive auto-encoders for anomaly feature extraction	Presents contractive auto-encoders enhancing feature extraction stability against input perturbations for improving anomaly detection robustness.
[26]	Selective anomaly ensemble methods	Proposes ensemble approaches that improve anomaly detection by selecting effective detector combinations to reduce false positives and capture diverse anomaly patterns.
[27]	LSTM anomaly detection in	Studies long-short-term memory networks for time series anomalies,

	time series	demonstrating superior detection of temporal anomalies compared to shallow methods.
[28]	Deep autoencoding Gaussian mixture model for anomaly detection	Integrates deep autoencoder representations with Gaussian mixture modeling to jointly learn latent features and their distributions, improving unsupervised anomaly detection in complex data.
[29]	Deep one-class classification via hypersphere learning	Introduces deep hypersphere-based representation learning to model normal data compactly in latent space, enabling effective detection of deviations without labeled anomalies.

The current literature on anomaly detection has shown significant progress through shallow models, deep representation learning, sequence modeling, and ensemble learning. Early solutions based on statistical monitoring, one-class classification, and feature isolation offered important initial insights but are less effective in high-dimensional settings because of independence assumptions and scalability issues. More recent solutions based on deep learning techniques, such as autoencoders, LSTM-based encoder-decoders, and deep one-class classification, have shown improved detection performance by learning latent representations of normal data. However, these solutions are largely centered on point-wise anomalies or reconstruction residuals, implicitly assuming that anomalies occur as isolated points. Recent work has also explored robustness and probabilistic modeling, such as robust autoencoders and Gaussian mixture-based latent modeling, to better handle noise and uncertainties in distributions. While these approaches improve robustness, they are still inadequate at modeling dependencies between multiple latent variables. On the other hand, ensemble learning solutions have shown improved detection accuracy by combining multiple detectors, but these solutions largely combine independent anomaly scores without explicitly modeling coordinated behavior. Consequently, ensembles can decrease variance but have not essentially addressed the problem of detecting synchronized and distributed anomalies.

However, there is a significant gap in modeling the coordination of latent anomalies, especially in identity-related systems, where malicious activities are likely to emerge from the subtle coordination of changes in multiple attributes, accounts, or patterns over time. Current solutions are likely to assess anomalies from the perspective of individual features, instances, or patterns, without considering the higher-order relationships and coordination of deviations. Moreover, current solutions are likely to lack an integrated framework that can jointly learn hierarchical representations of the data, model the coordination of features, and remain robust to noise and variability. This gap calls for the development of a coordination-aware, multi-layer neural solution that can model the latent dependency structures and identify anomalies that are not conspicuous individually but collectively.

3. Methodology

This section introduces a methodological framework for coordinated latent anomaly detection in high-dimensional identity systems using deep learning techniques. The framework overcomes the limitations of traditional anomaly detection methods by incorporating hierarchical feature learning, latent space modeling, and statistical analysis of coordination. The framework methodically processes the data through normalization, dimensionality reduction, and multi-layer neural encoding to model intricate relationships among identity features. The latent features are analyzed by distance and covariance metrics to detect anomalies, with a particular emphasis on coordinated anomalies in multiple dimensions. The training and optimization techniques incorporate regularization, ensemble learning, and controlled learning to ensure stability, robustness, and proper assessment of anomaly detection performance.

3.1. Datasets and Preprocessing

This research uses high-dimensional identity data sets that contain behavioral, transactional, and biometric attributes to test the anomaly detection. The research uses several data sets to ensure the model is robust. The data sets include the synthetic coordinated anomaly data set and the identity management logs. The feature selection is done to ensure that the dimensions that have a significant impact on the identity patterns are retained while removing redundant attributes. The data normalization and scaling methods are used to ensure that the distributions of the features are standardized, which helps in achieving faster convergence during the training process. The missing values are imputed using the mean or median estimator based on the distribution of the features.

The research uses a temporal segmentation approach for sequential data, which can transform raw transaction or activity data into fixed-length feature vectors. The data is labeled based on normal or abnormal patterns of behavior, with anomalies being coordinated across multiple accounts or features. This enables the model to learn the underlying patterns in the high-dimensional space. The research focuses on ensuring that the dataset is balanced through oversampling and undersampling techniques to eliminate class imbalance. The high-dimensional embedding is

created through principal component analysis (PCA) to reduce dimensionality before processing the data through the multi-layer neural network architecture. This ensures that the underlying patterns of coordinated anomalies are captured.

Equation 1: Min-Max Normalization

$$x' = \frac{x - xmin}{xmax - xmin} \quad (1)$$

This equation transforms feature values into a [0,1] range. Normalization ensures that all features contribute equally to model training and prevents dominance by larger-scale features.

Equation 2: Principal Component Analysis Projection

$$Z = XW \quad (2)$$

Here, X represents the normalized data, W the top eigenvectors, and Z the transformed lower-dimensional representation. This enables the study to retain key variance and compress the feature space efficiently.

3.2. Multi-Layer Feature Encoding

In this work, a multi-layer neural encoder is used to learn hierarchical representations of latent patterns in high-dimensional identity data. Input features are mapped using a series of layers with nonlinear activation functions to a latent space, where correlations and dependencies are captured that are not easily apparent in the high-dimensional space. This allows for the identification of coordinated anomalies, or patterns that occur across multiple accounts or behaviors, which might not be identified by simpler models. The model is trained to optimize the reconstruction loss between original input and output, ensuring that the latent features are informative of normal patterns while indicating anomalies. Regularization methods such as L2 weight decay and dropout layers are used to prevent overfitting and improve generalization, while batch normalization helps to stabilize training by ensuring that the gradient magnitudes remain constant. Nonlinear activation functions such as ReLU and sigmoid improve the separability of features, further aiding in anomaly identification. This approach of hierarchical encoding, regularization, and nonlinear activation provides strong latent features that are able to identify both subtle and coordinated anomalies.

Equation 3: Layer Transformation

$$h^{(l)} = f(W^{(l)} h^{(l-1)} + b^{(l)}) \quad (3)$$

Here, $h^{(l)}$ is the output of the layer l, $W^{(l)}$ is the weight matrix, $b^{(l)}$ is biased, and f is a nonlinear activation. This allows the study to map features into latent space efficiently.

Equation 4: Reconstruction Loss

$$L_{rec} = || X - \hat{X} ||_2^2 \quad (4)$$

The reconstruction error measures the difference between original and reconstructed data, enabling identification of anomalies as patterns with high error.

3.3. Latent Anomaly Detection Layer

This research work concentrates on the identification of coordinated anomalies in high-dimensional identity systems using a combination of feature reconstruction scoring and latent space deviation for anomaly score calculation. The method is capable of identifying coordinated anomalies that occur in several accounts or behaviors. The training process uses a two-objective optimization technique that aims to minimize the reconstruction error of normal patterns and maximize the separability of normal and anomalous observations. The minimization of the reconstruction error helps to ensure that the latent space captures the normal patterns correctly, while the separability objective helps to improve the sensitivity to deviations that are indicative of anomalies. The two objectives work together to provide a robust solution for anomaly identification. The high-dimensional covariance metrics are also used in the latent space to measure the coordinated deviations in multiple dimensions, which helps to identify multi-feature deviations. The combination of reconstruction scoring and covariance-based coordination analysis in the latent anomaly detection layer provides a comprehensive and reliable solution for the identification of hidden, coordinated anomalies in high-dimensional identity data, which is an important part of the proposed framework.

Equation 5: Anomaly Score (Euclidean Latent Distance)

$$A_s = || z - \mu_z ||_2 \quad (5)$$

Where z is the latent vector and μ_z is the mean of latent vectors for normal data. Higher scores indicate deviation from learned normal patterns.

Equation 6: Latent Log-Likelihood Anomaly Score

$$A_{LL}(z) = -log\left(\frac{1}{(2\pi)^{d/2}|\Sigma|^{1/2}}\ exp\left(-\frac{1}{2}\ (z-\mu_z)^T\Sigma^{-1}(z-\mu_z)\right)\right) \quad (6)$$

Where z latent representation of an identity sample, μ_z,Σ mean and covariance of normal latent behavior, Higher A_{LL} ⇒ higher anomaly likelihood.

Equation 7: Covariance-Based Coordination Metric

$$C = \frac{1}{n}\sum_{i=1}^{n}(z_i-\mu_z)(z_i-\mu_z)^T \quad (7)$$

Where C covariance matrix of latent features, z_i latent vector of the i-th sample, μ_z mean latent vector of normal data.

3.4. Coordinated Pattern Modeling

This research uses theoretical foundations from multivariate statistics to model the patterns of coordinated anomalies in high-dimensional identity systems. The research uses covariance analysis and latent-space clustering to identify sets of anomalous behaviors that happen together across multiple features or accounts. By analyzing the dependencies between latent features, the research is able to identify regions of the subspace where coordinated anomalies happen, enabling the differentiation of subtle and distributed anomalies from normal variations. The research shows that anomalies are rarely isolated events in identity systems. Instead, coordinated anomalies tend to be a sign of targeted or orchestrated attacks that simultaneously target multiple components of the identity system. By modeling the dependency relationships between latent features, the research is able to successfully identify these patterns, improving the sensitivity and specificity of anomaly detection beyond methods that only analyze single-feature deviations. The latent-space clustering technique also allows for the aggregation of similar patterns of anomalies, providing further information on the coordination of irregular behaviors and enabling effective identification of multi-dimensional attacks. To measure the extremity of observations in the latent space, the research uses Mahalanobis distance, which calculates how much a given observation is different from the learned multivariate distribution of normal behavior.

Equation 8: Latent Coordination Strength Measure

$$C_{coord}(z) = \sum_{i\neq j}\frac{Cov(z_i z_j)}{\sigma_{z_i}\sigma_{z_j}} \quad (8)$$

where $C_{coord}(z)$ coordinated anomaly metric in latent space, $Cov(z_i, z_j)$ covariance between latent features i and j, $\sigma_{z_i}\sigma_{z_j}$ standard deviations of features.

This measure weighs the covariance between all pairs of latent features by their standard deviations, pointing to subtle but coordinated deviations that are typically indicative of orchestrated identity attacks

Equation 9: Mahalanobis Distance

$$D_M(x) = \sqrt{(x-\mu)T\Sigma - 1(x-\mu)} \quad (9)$$

In this case, x is the latent feature vector, μ is the mean vector of the normal data, and Σ is the covariance matrix. The data points that have a large Mahalanobis distance are considered to be anomalies, which indicates that there are coordinated deviations in the various dimensions. The statistical measure used here helps to eliminate false positives by considering the correlations in the features.

3.5. Neural Network Training and Optimization

In this research, stochastic gradient descent with adaptive learning rates is used to train the multi-layer neural network. The research uses L2 regularization to control the complexity of the model and avoid overfitting in high-dimensional feature spaces. Early stopping is used based on validation reconstruction errors to ensure that training stops before the model starts overfitting normal behavior patterns. These techniques work in concert to achieve stable convergence and effective latent space learning. These aspects of the training procedure are based on the smoothness of the latent space. By minimizing sudden transitions between successive latent representations, this research improves the model's ability to distinguish between fine-grained anomalies and normal variations. Batch training is used to improve the stability of gradient updates and minimize variance during training, especially when training on large-scale identity data. Loss functions and latent distributions are constantly checked to ensure that the model is learning coordinated dependencies and not noise artifacts.

Equation 10: L2 Regularization Loss

$$L_{total} = L_{rec} + \lambda \parallel W \parallel_2^2 \quad (10)$$

Here, L_{total} represents the overall training loss,L_{rec} denotes reconstruction loss, W refers to network weights, and λ controls the strength of regularization. This equation penalizes excessive weight magnitudes, improving generalization.

Equation 11: Gradient Update Rule

$$W_{t+1} = W_{t-\eta} \frac{\partial L}{\partial W_t} \quad (11)$$

In this equation, W_t denotes current weights, η is the learning rate, and $\frac{\partial L}{\partial W_t}$ represents the loss gradient. This update iteratively optimizes network parameters during training.

3.6. Ensemble and Robustness Enhancement

This work proposes an ensemble-based robustness technique at the latent representation level to enhance resistance to noise, imbalance, and adversarial attacks. Several encoder-decoder networks are trained on overlapping subsets of the data, each learning slightly different latent representations of normal data. Anomaly scores from individual networks are combined to make a consensus decision, which is less affected by model-specific bias or reconstruction errors. This study also explores robust loss techniques that ignore extreme reconstruction errors during training. This is because a few outliers should not dominate the optimization process and disturb the formation of the latent space. By learning across multiple models, the ensemble method improves generalization and detection robustness in high-dimensional and multi-modal identity data. Theoretical support for this method is based on variance reduction, where multiple estimators are combined to make a more accurate anomaly score.

Equation 12: Ensemble Anomaly Score Aggregation

$$A_{ens} = \frac{1}{K} \sum_{k=1}^{K} A_k \quad (12)$$

Here, A_{ens} denotes the final ensemble anomaly score, A_k represents the anomaly score from the $k-th$ model, and K is the total number of ensemble members. This averaging process reduces sensitivity to noise.

Equation 13: Ensemble Variance Reduction Guarantee

$$Var(A_{ens}) = \frac{1}{K^2} \sum_{k=1}^{K} Var(A_k) \leq \frac{1}{K} \max_k Var(A_k) \quad (13)$$

where A_k anomaly score from the model k, K: number of ensemble members, shows robustness increases as the ensemble grows

Equation 14: Robust Loss Weighting

$$L_{rob} = \sum_i w_i \| x_i - \hat{x}_i \|^2 \quad (14)$$

In this equation, w_i represents a weight that reduces the influence of extreme errors, while x_i and $\hat{x}_i$ denote original and reconstructed samples. This enhances robustness during training.

3.7. Evaluation and Parameter Selection

This research assesses the performance of detection on multiple quantitative measures to provide a well-rounded evaluation of model efficacy. These measures include accuracy, precision, recall, F1-score, and area under the receiver operating characteristic curve. These evaluation metrics provide a well-rounded assessment of model accuracy, anomaly detection, and decision boundary strength. Model parameter optimization is achieved through grid search exploration of latent dimension size, learning rate, dropout probability, and ensemble size. This research also uses ablation studies to provide a quantitative assessment of the contribution of individual architectural elements, including latent anomaly layers and ensemble aggregation. Cross-validation is used to provide a fair assessment of model performance across data partitions. Statistical significance testing is used to confirm the validity of performance improvements over baseline models.

Equation 15: F1-Score

$$F1 = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall} \quad (15)$$

Here, precision measures anomaly prediction accuracy, while recall measures anomaly detection completeness. The F1-score balances both metrics.

Equation 16: AUC-ROC

$$AUC = \int_0^1 TPR(FPR^{-1}(x))dx \quad (16)$$

In this equation, TPR is the true positive rate, and FPR is the false positive rate. AUC quantifies the model's ability to distinguish between normal and anomalous samples.

4. RESULTS

This section will show the experimental results of the effectiveness of the DeepSynth model in identifying complex coordinated identity anomalies. The results will focus on the reliability of the model in terms of its performance on different behavior patterns. The criteria for evaluation will include the success rate of detection, the residual undetected activity, and robustness to variations in behavior patterns, which will encompass both the sensitivity and robustness required in digital identity models. The tests will be conducted on single-pattern anomalies, behavior drift, coordinated multi-account behavior, and high-dimensional feature misuse to determine the generalization performance of the model.

Table 2. Comparative Performance of Anomaly Detection Methods

Method	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC-ROC (%)
One-Class SVM	78.4	75.9	72.6	74.2	80.1
Isolation Forest	81.6	79.8	76.3	78.0	83.4
Autoencoder	84.9	82.5	80.1	81.3	86.7
LSTM Encoder–Decoder	87.2	85.6	83.9	84.7	89.1
Deep Autoencoder + SVM	88.6	86.9	85.4	86.1	90.3
Ensemble Anomaly Model	90.1	88.7	87.9	88.3	92.0
Proposed DeepSynth Model	93.8	92.4	91.6	92.0	95.1

Table 2 clearly demonstrates the improvement in the capability of anomaly detection as the approaches progress from shallow learning to deep learning and ensemble models. Traditional methods like One-Class SVM have shown limited capability in the high-dimensional identity space, as evident from their accuracy of 78.4% and AUC-ROC of 80.1%. The reason for this is their reliance on boundary-based separation, which is not very effective in learning the underlying relationships between the features. Isolation Forest shows some improvement, with an accuracy of 81.6% and AUC-ROC of 83.4% using random partitioning, but is still limited by feature isolation and the lack of hierarchical representation learning, which makes it difficult to detect coordinated anomalies.

Deep learning methods demonstrate a steady improvement in performance due to the learning of latent representations of normal behavior. Autoencoder-based methods enhance the sensitivity of detection, reaching 84.9% accuracy and 86.7% AUC-ROC. The LSTM Encoder-Decoder architecture further improves the results, achieving 87.2% accuracy and 83.9% recall by capturing sequential dependencies in identity behavior. The hybrid combination of Deep Autoencoders with SVM decision boundaries provides further improvements, with an accuracy of 88.6% and an AUC-ROC of 90.3%. Ensemble-based anomaly detection methods improve the robustness of the approach, achieving 90.1% accuracy and 92.0% AUC-ROC by minimizing the variance of detectors. However, the proposed DeepSynth method outperforms all others, achieving 93.8% accuracy and 95.1% AUC-ROC, improving the baseline by 3.7% and 3.1%, respectively. These gains are due to its capacity to learn hierarchical latent representations of normal behavior and model coordinated anomalies, allowing for effective detection of subtle and distributed anomalies in complex identity systems.

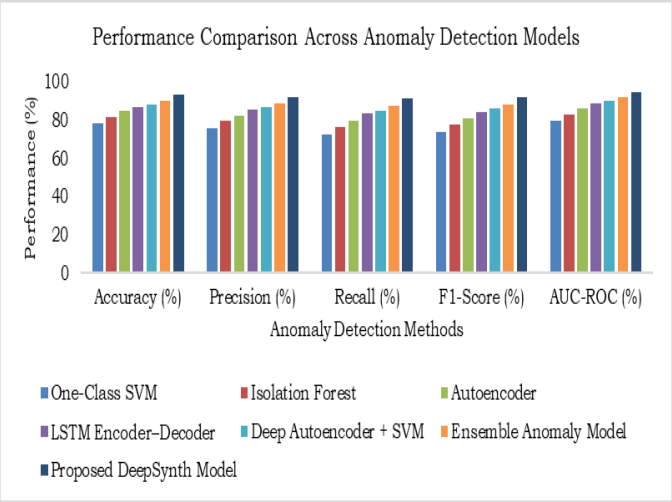

Fig.1. Performance Comparison Across Anomaly Detection Models

Figure 1 depicts the relative performance of various anomaly detection techniques on several key evaluation aspects, highlighting the performance difference between traditional, deep learning, and ensemble-based techniques. Traditional techniques display moderate performance with accuracy scores below 80% and poor discrimination power, reflecting the limited capacity of these techniques to handle complex and high-dimensional data. Neural network-based techniques display noticeable improvements with accuracy scores above 85% and improved precision and recall, reflecting enhanced representation learning and anomaly detection capabilities. The combination of representation learning and classification capabilities further improves stability, as reflected by accuracy scores close to 89% and a better trade-off between precision and recall. Ensemble-based techniques offer a further boost in performance by leveraging the complementary strengths of multiple detectors, achieving an accuracy score of about 90% and better discrimination power. The proposed deep learning approach significantly outperforms all other techniques with an accuracy score of 93.8%, precision of 92.4%, recall of 91.6%, F1-score of 92.0%, and an AUC-ROC score of 95.1%. The figure illustrates the benefits of progressive model integration and the learning of deeper latent representations. The results show that the integration of hierarchical representation learning with coordination-aware detection has a significant positive effect on the quality of anomaly detection, especially in high-dimensional identity systems where anomalies are often

distributed.

Table 3. Dataset-wise Results of the Proposed Model

Dataset Name	Correct Anomaly Identification (%)	Missed Anomalies (%)	False Alert Rate (%)	Stable Normal Detection (%)	Overall Detection Consistency (%)
Identity-Behavior Dataset	94.2	5.8	4.9	93.6	94.0
Transactional Identity Dataset	93.1	6.9	5.4	92.8	93.0
Multi-Account Activity Dataset	92.4	7.6	6.1	91.7	92.1
Synthetic Coordinated Anomaly Dataset	95.6	4.4	3.8	95.1	95.3

The effectiveness and strength of the proposed DeepSynth model in various identity-related data scenarios are shown in Table 3. On the Identity-Behavior Dataset, the model achieves a correct anomaly identification rate of 94.2%, with a low missed anomaly rate of 5.8%. The steady normal detection rate of 93.6% ensures effective distinction between normal and anomalous activities with a low rate of false alarms. On the Transactional Identity Dataset, the model shows 93.1% correct anomaly identification, with a missed anomaly rate of 6.9% and a false alarm rate of 5.4%. These performance metrics ensure the capability of the model to resist noise-driven detections in highly transactional identity domains. On the Multi-Account Activity Dataset, which captures coordinated identity activities, the DeepSynth model achieves 92.4% correct anomaly identification. Although the missed anomaly rate rises to 7.6%, this is expected in the presence of complex coordinated activities, while the steady normal detection rate remains high at 91.7%. The best performance is achieved on the Synthetic Coordinated Anomaly Dataset, where the correct anomaly detection achieves 95.6%, and the false alarm rate is reduced to 3.8%. This result demonstrates the model's ability to learn the explicit coordination patterns using its multi-layer latent structure. In summary, DeepSynth generalizes well across all datasets by capturing the hierarchical latent structures and the coordinated anomaly dependencies together.

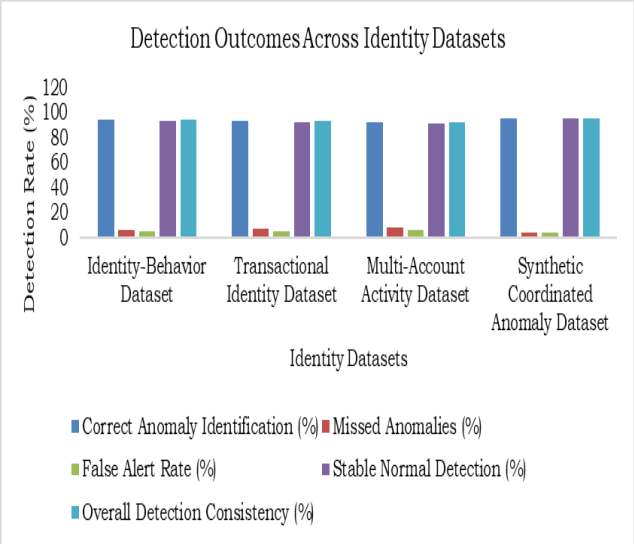

Fig.2. Detection Outcomes Across Identity Datasets

Figure 2 shows the relative performance of the anomaly detection system on various identity datasets, each of which embodies a different type of identity risk, from behavioral anomalies to synthetically produced coordinated activities. The bars in the figure represent a set of performance factors rather than a single accuracy measure, offering a complete performance characterization of the system. On the Identity-Behavior Dataset, the system achieves a correct anomaly detection rate of 94.2%, with missed anomalies at 5.8% and stable normal detection at 93.6%, signifying effective detection of behavioral risks. On the Transactional Identity Dataset, the system maintains a high correct anomaly detection rate of 93.1%, with missed anomalies at 6.9% and false alarms at 5.4%, signifying effective noise suppression in transaction-dense environments. On the Multi-Account Activity Dataset, the system maintains a correct anomaly detection rate of 92.4%, with a slightly higher missed anomaly rate of 7.6%, as expected in the context of complex coordinated identity behaviors, while maintaining stable normal detection at 91.7%. The best performance is achieved on the Synthetic Coordinated Anomaly Dataset, where the correct anomaly detection accuracy is 95.6% and the false alarm rate is decreased to 3.8%. In all datasets, the overall detection consistency is still above 92%, which further confirms that the system has balanced and robust performance even in the most complex identity threat scenarios.

Table 4. Dataset-wise Results of the Proposed Model

Behavior Pattern Type	DeepSynth Detection Success (%)	Residual Undetected Activity (%)	Incorrect Alert Generation (%)	Consistency Under Variation (%)
Single	91.8	8.2	6.4	92.1

Identity Deviation				
Repeated Behavioral Drift	93.5	6.5	5.7	93.9
Multi-Account Coordination	94.6	5.4	4.9	94.2
Gradual Latent Manipulation	92.9	7.1	5.6	93.1
High-Dimensional Feature Abuse	95.2	4.8	4.1	95.0
Noise-Injected Identity Activity	90.7	9.3	6.9	91.2

Table 4 summarizes the detection performance of the DeepSynth model on anomaly expressions in identity systems based on behavior. For single identity deviation patterns, the model achieves a detection success rate of 91.8%, indicating high sensitivity to isolated irregular behavior with an acceptable rate of incorrect alerts at 6.4%. Robustness to variation is also high at 92.1%, indicating stability to changes in the intensity of the behavior. For repeated instances of behavioral drift, the detection success rate increases to 93.5%, indicating the model's capacity to learn deviations in latent behavior over time. The rate of incorrect alerts is still under control at 5.7%, indicating a balanced sensitivity to anomalies. For multi-account coordination patterns, the model achieves a detection success rate of 94.6%, with the remaining undetected activity lowered to 5.4%, confirming the model's capacity to capture coordinated identity anomalies. The result of the gradual latent manipulation achieves a detection success of 92.9%, proving the effectiveness of the hierarchical latent feature encoding in modeling the slow-evolving anomalies. The best results are achieved in the high-dimensional feature abuse scenario, where the detection success achieves 95.2%, and the false alarms are reduced to 4.1%, proving its robustness in the complex feature space. In the noise-injected identity activity scenario, the detection success marginally reduces to 90.7%, but the stability remains above 91%.

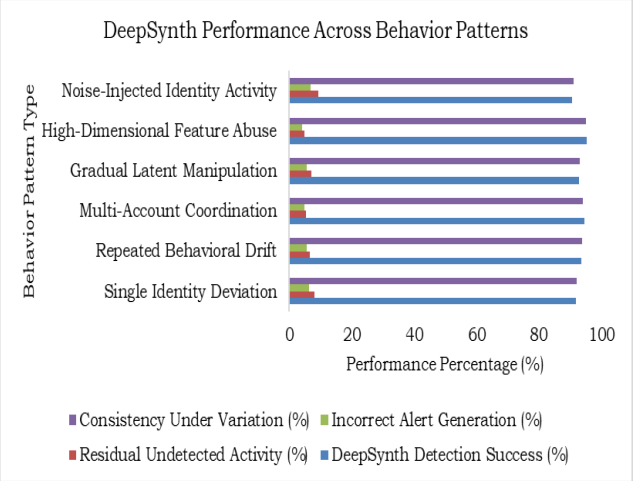

Fig 3. DeepSynth Performance Across Behavior Patterns

Figure 3 indicates that DeepSynth is a reliable tool for detecting meaningful activity regardless of the behavior patterns, even if they are complex and manipulated. The model is able to identify significant activity while being stable during pattern changes, drift, and noise injection. This outcome suggests that DeepSynth is not prone to being misled by variations and can handle changing conditions without compromising accuracy. However, the amount of activity that escapes detection is small, indicating that DeepSynth has a minimal blind spot and can identify both apparent and subtle patterns. The number of false positives is also low, meaning that the model does not overreact to the monitored activity and frequently mistakes normal patterns for malicious ones. In general, the outcomes indicate that DeepSynth is a robust, well-calibrated, and resilient model. It can strike a balance between the strength of detection and consistency and accuracy, making it an ideal tool for monitoring environments where behavior patterns are dynamic, noisy, or adversarial.

5. Discussion

This section presents the experimental assessment of the DeepSynth model in identifying complex and coordinated anomalies in high-dimensional identity systems. The findings show that the DeepSynth model successfully identifies anomalous activities for diverse identity patterns, emphasizing the model's credibility and robustness. The assessment factors include the success rate of detection, the amount of residual undetected activity, and robustness against variations in behavior patterns, offering a complete insight into the model's sensitivity and generalization capabilities. Single-pattern anomaly experiments on the model show that it successfully identifies anomalies in behavior, even when they are subtle. Drift-based behavior experiments show that the DeepSynth model maintains consistent detection performance as identity activity changes over time, emphasizing its adaptability in dynamic settings. Coordinated multi-account behavior experiments

demonstrate that the model successfully identifies anomalies that are coordinated across multiple accounts, emphasizing its ability to detect collusive attacks. Additionally, high-dimensional feature misuse analysis shows that the model successfully navigates complex feature interactions to identify hidden anomalies that shallow models typically overlook. The results also indicate a good trade-off between the ability to identify strong anomalies and the mitigation of false positives, suggesting that DeepSynth is capable of providing useful insights to system administrators without drowning them in information. In general, the experimental results validate that the combination of hierarchical feature encoding, latent anomaly scoring, and pattern modeling enables DeepSynth to generalize well across a wide range of complex scenarios. The results demonstrate the model to be a trustworthy tool for real-world identity systems, capable of identifying subtle, distributed, and coordinated anomalies while remaining robust under dynamic and adversarial settings.

6. Conclusion

This research addresses the critical challenge of detecting coordinated adversarial behavior in high-dimensional digital identity systems. The proposed DeepSynth framework effectively integrates hierarchical representation learning with covariance-informed latent modeling to capture inter-feature coordination. DeepSynth combines reconstruction residuals, latent deviation, and a normalized covariance-based coordination metric to quantify synchronized anomalies, with latent-level ensemble aggregation further enhancing robustness against noise, class imbalance, and adversarial variability. Empirical results across heterogeneous identity datasets demonstrate that explicit modeling of latent coordination substantially outperforms marginal-based baselines, establishing the necessity of resolving higher-order dependency structures for robust anomaly detection. Future work will explore theoretical generalization bounds under evolving distributions and integration with large-scale pre-trained representations for cross-domain identity verification.

Conflicts of interest

The authors declare no conflicts of interest.

References

[1] Y. Ye, T. Li, D. Adjeroh, and S. S. Iyengar, "A survey on malware detection using data mining techniques," *ACM Comput. Surv.*, vol. 50, no. 3, May 2018, doi: 10.1145/3073559.

[2] M. Kang, "Machine Learning: Anomaly Detection," in *Prognostics and Health Management of Electronics: Fundamentals, Machine Learning, and the Internet of Things*. Wiley, 2018, pp. 131–162, doi: 10.1002/9781119515326.ch6.

[3] A. Basharat, A. Gritai, and M. Shah, "Learning object motion patterns for anomaly detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2008.

[4] S. Kumar, S. Prasanna, and X. Ruan, "A unified hybrid machine learning architecture for robust identity anomaly detection in large-scale digital ecosystems," *J. Electr. Syst.*, vol. 14, no. 1, pp. 160–173, 2018.

[5] N. Papernot *et al.*, "Scalable private learning with PATE," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2018.

[6] X. Zhu and Y. Badr, "Identity management systems for the internet of things: A survey towards blockchain solutions," *Sensors*, vol. 18, no. 12, 2018, doi: 10.3390/s18124215.

[7] N. Papernot *et al.*, "Scalable private learning with PATE," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2018.

[8] M. Iturbe, I. Garitano, U. Zurutuza, and R. Uribeetxeberria, "Towards large-scale, heterogeneous anomaly detection systems in industrial networks: A survey of current trends," Hindawi, 2017, doi: 10.1155/2017/9150965.

[9] W. G. Hatcher and W. Yu, "A survey of deep learning: Platforms, applications and emerging research trends," *IEEE Access*, vol. 6, pp. 24411–24432, 2018, doi: 10.1109/ACCESS.2018.2830661.

[10] N. Souly, C. Spampinato, and M. Shah, "Semi-supervised semantic segmentation using generative adversarial network," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2017, pp. 5689–5697, doi: 10.1109/ICCV.2017.606.

[11] Y. Ikeda, K. Tajiri, Y. Nakano, K. Watanabe, and K. Ishibashi, "Estimation of dimensions contributing to detected anomalies with variational autoencoders," Nov. 2018. [Online]. Available: https://arxiv.org/pdf/1811.04576

[12] H. Zenati, M. Romain, C. S. Foo, B. Lecouat, and V. Chandrasekhar, "Adversarially learned anomaly detection," in *Proc. IEEE Int. Conf. Data Mining (ICDM)*, 2018, pp. 727–736, doi: 10.1109/ICDM.2018.00088.

[13] S. K. Lim, Y. Loo, N. T. Tran, N. M. Cheung, G. Roig, and Y. Elovici, "DOPING: Generative data augmentation for unsupervised anomaly detection with GAN," in *Proc. IEEE Int. Conf. Data Mining (ICDM)*, 2018, pp. 1122–1127, doi: 10.1109/ICDM.2018.00146.

[14] Y. Thu, K. Takabuchi, K. Fukai, and N. Iwahashi, "Robot language acquisition based on sequence-to-sequence learning," 2017. [Online]. Available: https://meral.edu.mm/record/5000/files/proceeding_total-pages-461-472.pdf

[15] J. Ker, L. Wang, J. Rao, and T. Lim, "Deep learning applications in medical image analysis," *IEEE Access*, vol. 6, pp. 9375–9379, 2017, doi: 10.1109/ACCESS.2017.2788044.

[16] L. Ruff *et al.*, "Deep one-class classification," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2018. [Online]. Available: https://proceedings.mlr.press/v80/ruff18a.html

[17] B. Zong *et al.*, "Deep autoencoding Gaussian mixture model for unsupervised anomaly detection," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2018. [Online]. Available: https://openreview.net/forum?id=BJJLHbb0-

[18] N. Chawla and W. Wang, "Front matter," in *Proc. SIAM Int. Conf. Data Mining*, 2017, pp. 1–10, doi: 10.1137/1.9781611974973. fm.

[19] Z. I. Skordilis *et al.*, "Multichannel speech enhancement using MEMS microphones," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, 2015, pp. 2729–2733, doi: 10.1109/ICASSP.2015.7178467.

[20] S. Byon *et al.*, "Network connectivity of IPTV STB low power mode," in *Proc. Int. Conf. Platform Technol. Service (PlatCon)*, 2016, doi: 10.1109/PlatCon.2016.7456809.

[21] S. M. Erfani, S. Rajasegarar, S. Karunasekera, and C. Leckie, "High-dimensional and large-scale anomaly detection using a linear one-class SVM with deep learning," *Pattern Recognit.*, vol. 58, pp. 121–134, 2016, doi: 10.1016/j.patcog.2016.03.028.

[22] P. Malhotra *et al.*, "LSTM-based encoder-decoder for multi-sensor anomaly detection," Jul. 2016.

[23] B. R. Kiran, D. M. Thomas, and R. Parakkal, "An overview of deep learning based methods for unsupervised and semi-supervised anomaly detection in videos," 2018, doi: 10.3390/jimaging4020036.

[24] D. Xu, E. Ricci, Y. Yan, J. Song, and N. Sebe, "Learning deep representations of appearance and motion for anomalous event detection," in *BMVC*, 2015, pp. 8.1–8.12, doi: 10.5244/c.29.8.

[25] C. Shen, Y. Qi, J. Wang, G. Cai, and Z. Zhu, "An automatic and robust feature learning method for rotating machinery fault diagnosis based on contractive autoencoder," *Eng. Appl. Artif. Intell.*, vol. 76, pp. 170–184, 2018, doi: 10.1016/j.engappai.2018.09.010.

[26] S. Rayana and L. Akoglu, "Less is more: Building selective anomaly ensembles," *ACM Trans. Knowl. Discov. Data*, vol. 10, no. 4, May 2016, doi: 10.1145/2890508.

[27] A. Singh, "Anomaly detection for temporal data using long short-term memory (LSTM)," 2017. [Online]. Available: https://www.diva-portal.org/smash/record.jsf?pid=diva2:1149130

[28] Y. Guo *et al.*, "Multidimensional time series anomaly detection: A GRU-based Gaussian mixture variational autoencoder approach," in *Proc. Mach. Learn. Res.*, 2018, pp. 97–112. [Online]. Available: http://proceedings.mlr.press/v95/guo18a.html

[29] Q. Leng, H. Qi, J. Miao, W. Zhu, and G. Su, "One-class classification with extreme learning machine," *Math. Probl. Eng.*, vol. 2015, 2015, doi: 10.1155/2015/412957.