10.48047/jocaaa.2022.30.02.66

# EvoAnom: Longitudinal Temporal Deep Learning for Evolutionary Anomaly Detection in Digital Identities

**Suman Kumar Sanjeev Prasanna[1*], Shardul Pandya[1]**
[1]School of Computer and Information Sciences, University of the Cumberlands, Williamsburg, KY
*Corresponding Author Email: sprasanna68498@ucumberlands.edu

**Abstract**

Identity fraud and behavioral anomalies evolve continuously, often rendering static detection methods ineffective. This research introduces EvoAnom, a longitudinal temporal deep learning framework for evolutionary anomaly detection in digital identity systems. By modeling sequential interactions as temporal embeddings, the framework captures both immediate irregularities and long-term behavioral shifts. Recurrent neural architectures, augmented with attention mechanisms and temporal convolutional layers, are employed to detect subtle deviations indicative of synthetic identity creation or account compromise. The study also incorporates temporal regularization and feature weighting to balance sensitivity to rare anomalies with stability in normal activity patterns. Empirical evaluation across multi-year, high-throughput identity datasets demonstrates that EvoAnom significantly outperforms baseline temporal and non-temporal models in anomaly detection accuracy, early threat identification, and long-term generalization. These findings highlight the value of temporal deep learning for continuous identity monitoring and adaptive security.

**Keywords:** Longitudinal modeling, Temporal deep learning, Evolutionary anomaly detection, Drift adaptation, Transfer learning, Sequential datasets, Robust representation.

## 1. INTRODUCTION

Anomaly detection has now become an integral part of intelligent monitoring systems. Modern digital infrastructures are characterized by data streams that are generated sequentially with evolving patterns over time, unlike static data [1]. Earlier, anomaly detection mechanisms were developed with the assumption that data is independent and identically distributed [2]. Such scenarios are rarely seen in real-world environments, where temporal dependencies and distributional changes are inherent features of data. With increasing complexity of systems, anomalies do not occur as isolated data points; they occur with gradual changes that are part of evolving temporal dependencies [3]. Deep learning has now been able to change the face of sequential data models by automatically learning high-dimensional representations of data. Recurrent neural networks and gated neural networks have shown promising results in learning temporal dependencies from structured sequences. These models are able to learn short-term and long-term dependencies by maintaining hidden states that learn from data sequences [4]. Convolutional and hybrid models of temporal neural networks are able to learn high-dimensional representations of data by learning local and hierarchical temporal dependencies. However, it is seen that most of the anomaly detection models still operate with static models, where models are learned and updated once.

One of the major issues in the analysis of longitudinal data is concept drift. Concept drift refers to changes in the statistical properties of the underlying data distributions. These changes can be caused by a variety of factors, including changes in user behavior, adversarial learning, seasonal changes, and updates [5]. The changes in the underlying data cause a problem because the learned models will start to diverge over time. The problem of concept drift is particularly troublesome in security-related domains because adversaries tend to keep changing their approaches in order to avoid being caught [6].

10.48047/jocaaa.2022.30.02.66

This causes evolutionary anomalies that are different in nature from other anomalies. Transfer learning has been found to be an effective solution for improving the performance of models in related domains. This reduces the dependence of models on a lot of labeled data. The use of transfer learning for handling time-related issues adds another layer of complexity [7]. The problem of changes in the underlying distributions of the data in the source and target domains has been addressed through the use of robust learning. The intersection of longitudinal modeling, temporal representation learning, and robustness analysis underscores the need for models that explicitly address evolutionary behavior. To detect anomalies in changing environments, models need to be able to track changes, adapt to drift, and maintain statistical reliability over long periods of time [8].

In this study, a strong cross-domain transfer learning approach is developed for pattern detection in changing digital identity systems. The major goal is to ensure accurate detection of fraudulent or anomalous behavior in heterogeneous systems, handling temporal changes and data scarcity issues. The scope includes several identity domain types, including financial, telecom, and e-commerce systems, where user behavior varies and is difficult to model using traditional supervised learning techniques. The problem justifies itself in today's environment, given the complexity of identity fraud, changing user behavior, and the necessity for effective generalization in heterogeneous systems. The goals of this study include developing a feature representation mechanism that aligns with the source and target domains, developing a temporal adaptation mechanism to account for temporal changes, and developing a robustness mechanism to resist adversarial attacks. This approach has shown promising results compared to existing approaches, with improved accuracy, temporal stability, and generalization performance. The structure of this paper includes an introduction to related work, problem definition, proposed approach, results, discussion, and conclusions.

## 2. LITERATURE SURVEY

Anomaly detection for sequential data and data that is heavily dependent on time has been an actively developing field of study, especially with the increasing availability of high-resolution time series data in domains like finance, cybersecurity, and IoT infrastructures. However, with the increasing complexity of time series data, conventional methods that use statistical thresholds or shallow machine learning models are proving to be inadequate. In this regard, deep learning models that account for time dependencies are effective alternatives that can learn long-term dependencies without requiring manual feature engineering. Long Short-Term Memory (LSTM) networks and its variants are effective for time-based anomaly detection. Various literature reviews highlight the importance of effective temporal representations, domain independence, and concept drift-aware models. This review focuses on foundational and advanced literature that influences our understanding of deep learning models for anomaly detection in time series data [9].

The research by Rongbin Xu et al. [10] discusses an improved LSTM-based anomaly detection method that takes into account an adaptive technique that can reduce the impact of concept drift on streaming IoT data. The research indicates that using an LSTM model that is normally trained under static conditions can perform poorly when there are statistical changes in the input streams, which is a common occurrence in smart cities. The research proposed a time-aware forgetting mechanism and a smooth activation function that can be integrated into an LSTM model for improved adaptability to concept drift and multi-class anomaly states. The research used real-world IoT communication datasets that showed improved accuracy, precision, and recall rates compared to conventional sampling methods, thus validating that concept drift can be improved by being integrated into an LSTM model. The research is important for understanding how deep sequential models can be improved for handling evolutionary changes that affect the quality of anomaly detection.

Shigeru Maya, Ken Ueno, and Takeichiro Nishikawa [11] propose a novel approach for anomaly

10.48047/jocaaa.2022.30.02.66

detection using a delayed LSTM (dLSTM), which uses multiple predictive models in an LSTM model to tackle issues like multi-mode normal behavior and prediction noise. In this approach, anomaly scores are determined based on prediction errors that are computed after receiving actual values, thus dynamically selecting models for improved accuracy in anomaly detection. This paper is significant as it highlights that anomalies in time sequences may not be easily determined using fixed windowing or a single prediction model; instead, dynamic prediction model selection based on actual outcomes makes it clearer to distinguish between normal and abnormal data sequences. This paper emphasizes the need for dynamic temporal predictions in deep learning models for anomaly detection in complex systems.

Alexander Geiger et al. [12] introduced an unsupervised framework based on GANs, named TadGAN, which uses LSTM networks along with cycle consistency reconstruction for quantifying anomalies. Adversarial learning is used for the generation of synthetic sequences that exhibit normal temporal correlation, while anomaly detection is achieved by checking the divergences between the reconstructed sequences by the generators and actual observations. This approach showed strong generalization capabilities over multiple public time-series datasets, achieving better performance than other baselines for F1 metrics. The contribution of this paper is the integration of generative learning with temporal modeling, showing that anomaly detection is achieved by the fidelity of reconstructed sequences as well as the quality of the distributional representations.

In the study by Tolga Ergen et al., [13] the authors examine the use of unsupervised and semi-supervised anomaly detection with LSTMs and one-class classification models like OC-SVM and SVDD. It is shown that by jointly learning effective representations of time series sequences and decision boundaries, it is possible to efficiently handle variable-length sequences without requiring any labeled abnormal data. It is emphasized that by effectively integrating temporal feature learning with effective classification boundaries, it is possible to achieve improved performance for time series data with limited or no available anomaly labels. Such models effectively bridge the gap between effective feature learning and statistical decision mechanisms, suggesting that deep models for temporal learning are useful feature extractors for subsequent decision mechanisms.

Lindemann et al. [14] A broader perspective on deep learning-based temporal anomaly detection methods is offered by the survey on LSTM anomaly detection, which discusses the variety of LSTM architecture variants and their appropriateness for dynamic context-dependent anomaly detection scenarios. This survey highlights the limitations of traditional statistical approaches in dealing with temporal and context-dependent dependencies in complex systems, while recurrent neural networks like LSTM allow for the capture of long-term patterns in sequential data. Another important aspect discussed in the survey is the emerging trends in LSTM-based anomaly detection, including graph-based and transfer learning approaches, which highlight the current research direction that goes beyond the basic LSTM architecture.

Table 1. Key Deep Learning Studies for Time-Series Anomaly Detection

| Study | Methods | Key Findings |
|---|---|---|
| [15] | Introduced an autocorrelation-based LSTM-Autoencoder for anomaly detection, leveraging temporal dependence and reconstruction error to flag anomalies in time series. | The study shows that incorporating autocorrelation features into LSTM-AE training improves detection accuracy compared to basic reconstruction approaches, especially in complex temporal sequences where pure sequence learning may miss structural dependencies. Model performance evaluation on several real-world datasets demonstrated robust anomaly identification for varied seasonal and trending data. |
| [16] | Developed a Temporal Convolutional Network Autoencoder (TCN-AE) | Results indicate that TCN-AE captures long-range temporal patterns more efficiently than traditional recurrent networks. The method achieved strong |

10.48047/jocaaa.2022.30.02.66

| | | |
|---|---|---|
| | using dilated causal convolutions to learn temporal representations for unsupervised detection. | unsupervised anomaly detection performance on electrocardiogram (ECG) benchmark data, highlighting effectiveness in capturing periodic and complex temporal trends without labeled anomalies. |
| [17] | Proposed Enhanced LSTM Autoencoders (ELSTMAE) with statistical filtering based on the central limit theorem to distinguish true anomalies from expected changes. | Findings show that using a theoretical probability criterion on residuals improves anomaly discrimination in streaming datasets. The enhanced training algorithm outperformed baseline LSTM approaches by better separating gradual and abrupt change signals within large unlabeled time series. |
| [18] | Explored active learning integrated with LSTM-autoencoder anomaly detection to improve detection with minimal labeled examples. | The study found that combining LSTM reconstructions with active annotation feedback enhanced model adaptability in semi-unsupervised environments. Active learning reduced false positives by refining boundary thresholds through expert-in-loop sample queries. |
| [19] | Presented a hybrid VAE-LSTM model for time series anomaly detection at an IEEE conference, combining probabilistic latent modeling with temporal sequence learning. | The model leveraged the strengths of variational encoding and LSTM sequence modeling to better capture temporal distribution and reconstruction uncertainty. Results on public benchmarks showed improved detection scores versus pure LSTM-AE and standalone statistical methods, with more reliable anomaly thresholds due to the probabilistic model's uncertainty estimation. |

The existing literature pertaining to the use of deep learning-based time series anomaly detection approaches has already proven its effectiveness in handling sequential dependencies. However, a number of structural issues remain open. The majority of existing approaches rely upon static frameworks and assume stationarity in the underlying data distributions. Real-world systems, however, tend to suffer from longitudinal drift caused by environmental changes, adversarial changes, and operational changes. Several approaches employing LSTM and autoencoder-based frameworks have been proposed for point anomaly detection. These approaches fail to capture the underlying drift in the time series. The use of transfer learning has been attempted in a few scenarios. However, the effectiveness of the approaches in handling cross-domain time series remains insufficient. The performance of the approaches has been evaluated in terms of accuracy without employing inferential statistics and analysis of detection delays. The proposed research aims to bridge the gap in the existing literature pertaining to time series anomaly detection. The proposed approach aims to incorporate a number of features in a unified optimization framework.

## 3. METHODOLOGY

This particular methodology explains a longitudinal temporal deep learning approach that is geared towards the detection of evolutionary anomalies within non-stationary data conditions. It should be noted that the approach to data modeling incorporates sequential data in a manner that retains temporal order while ensuring causality and behavioral evolution over time. As opposed to static approaches to anomaly detection, the proposed approach combines temporal representation learning, adaptation to drifts, transfer initialization, and robust optimization in one unified process. As such, the proposed approach focuses on statistically grounded learning, wherein descriptive and inferential evaluation approaches inform the process of parameter optimization and validation. Training within the proposed approach utilizes forward temporal validation to avoid data leakage and simulate actual deployment conditions. Moreover, the proposed approach incorporates an architecture that learns stable representations while remaining dynamically adapted to changes in the data distribution. As such, drift-

10.48047/jocaaa.2022.30.02.66

sensitive regularization promotes adaptability without leading to overfitting, while robust optimization enhances the overall robustness to noise and adversarial data.

### 3.1 Longitudinal Data Modeling and Preprocessing Framework

This process starts by developing a longitudinal data set representation that is appropriate for evolutionary anomaly detection. This data set consists of multivariate time series data collected over large periods of time, including structured transactional data, sensor data, and behavioral event data. This data is then divided into ordered temporal windows to preserve temporal integrity during training and testing. In this research, forward-chaining temporal data splitting is used to avoid information leakage during testing. Normalization of the data is done using window-specific statistical methods to preserve temporal integrity and avoid distribution contamination. Missing data values are handled using time-aware interpolation methods to avoid altering data continuity.

In this process, the training data is divided into rolling temporal windows to enable the data to learn short-term and long-term dependencies. It is assumed that there may be partially available anomaly labels during training. Therefore, a semi-supervised approach with dominant normal patterns during early stages of training is used. Feature scaling is done using longitudinal segments to handle gradual drifts in data.

Equation 1: Longitudinal Dataset Representation

$$D = \{(x_t, y_t)\}_{t=1}^{T} \tag{1}$$

Here, $x_t$ denotes the feature vector at time $t$, and $y_t$ denotes anomaly label.

Equation 2: Window Segmentation Function

$$W_i = \{x_t\}_{t=i}^{i+\tau} \tag{2}$$

Here, $\tau$ denotes window size and $i$ denotes starting index.

The dataset modeling stage ensures statistical coherence, preserves temporal causality, and prepares the input for adaptive training under evolutionary conditions.

### 3.2 Temporal Representation Learning Architecture

The current work proposes a gated recurrent temporal encoder. The proposed architecture aims at modeling dynamic dependencies in a sequence of segments. This architecture comprises a multi-layer recurrent unit that maintains memory states over a given window. The optimization of the proposed model aims at reconstructing and aligning the prediction.

The proposed research utilizes batch-wise temporal progression. The batch progression maintains a sequential ordering. The proposed model utilizes gradient backpropagation through time for iterative updates. Regularization has been added to prevent overfitting during sparse anomaly detection. The proposed model extracts latent representations of the evolving temporal patterns and filters out noise.

Equation 1: Hidden State Update

$$h_t = f(W_{xt} + Uh_{t-1}) \tag{3}$$

Here, $h_t$ denotes hidden state, $W$ and $U$ denote weight matrices

Equation 2: Reconstruction Loss

$$L_{rec} = || \, x_t - \hat{x}_t \, ||^2 \tag{4}$$

Here, $\hat{x}_t$ denotes reconstructed output.

This temporal representation stage enables adaptive memory retention and robust embedding generation across evolving sequences.

### 3.3 Drift-Aware Evolution Modeling

10.48047/jocaaa.2022.30.02.66

The research also introduces a drift-sensitive regularization mechanism used in the quantification of the evolution of the distribution in longitudinal windows. Unlike the traditional approach of assuming stationarity, the framework measures the divergence between successive latent distributions. If the divergence exceeds a learned threshold, adaptive retraining is initiated.

The training procedure also considers the incorporation of drift penalties in the objective function for the model's stability in the context of temporal change. Statistical consistency is also encouraged between successive segments.

Equation 1: Drift Score

$$S_t = || \mu_t - \mu_{t-1} || \tag{5}$$

Here, $\mu_t$ denotes the mean latent representation.

Equation 2: Drift Regularization Loss

$$L_{drift} = D_{KL}(P_t \,||\, P_{t-1}) \tag{6}$$

Here, $P_t$ denotes the latent distribution at time $t$.

This mechanism ensures evolutionary sensitivity while maintaining structural continuity in learned representations.

### 3.4 Transfer and Robustness Optimization

The current work incorporates the transfer-aware initialization technique to improve generalization with scarce labeled data. The parameters are learned from a source temporal domain and then fine-tuned for the target longitudinal segments, thereby reducing training instabilities and improving the accuracy of early detection results. Robustness is also improved through the application of perturbation training. During training, the model is exposed to small, controlled perturbations in the input sequence, representing adversarial or noisy conditions, to develop invariance to minor changes in the data distribution.

Equation 1: Transfer Update Rule

$$\theta' = \theta - \eta \nabla L \tag{7}$$

Here, θ denotes parameters and η denotes the learning rate.

Equation 2: Adversarial Perturbation

$$x^{adv} = x + \epsilon \cdot sign(\nabla_x L) \tag{8}$$

Here, ϵ denotes perturbation magnitude.

This optimization strategy enhances stability and improves detection reliability under dynamic environments.

### 3.5 Statistical Evaluation and Training Parameters

The final methodological step is the definition of the training parameters and the performance measures. The research uses mini-batch gradient descent with adaptive learning rate scheduling for training the model. The training process is done for a certain number of epochs with early stopping criteria for the validation loss function. The hyperparameters used in the research include the window size, hidden dimension, learning rate, drift threshold, and regularization coefficients. The hyperparameters are tuned through the grid search technique with temporal cross-validation.

The performance measures used to validate the research model include accuracy, precision, recall, F1 score, and area under the ROC curve. Detection delay is also used to validate the research model's responsiveness to new anomalies that may occur in the system. Statistical significance is also used to

10.48047/jocaaa.2022.30.02.66

validate the improvements in the research model compared to the baseline models through the paired sample t-test, and the confidence interval is used to estimate the uncertainty in the performance measures.

Equation 1: F1 Score

$$F1 = \frac{2PR}{P+R} \tag{9}$$

Here, $P$ denotes precision and $R$ denotes recall.

Equation 2: Detection Delay

$$Delay = t_{detect} - t_{actual} \tag{10}$$

Here, $t_{detect}$ denotes predicted anomaly time.

This evaluation framework ensures quantitative rigor and statistically reliable validation of the longitudinal temporal deep learning approach.

## 4. RESULTS

The results achieved in this study validate the efficacy of the proposed Longitudinal Temporal Deep Learning framework. It is clear from the results that the proposed model is effective for different sequential datasets. It is also important to note that the proposed model has been tested for different application domains, including network traffic, financial fraud, IoT sensor, and healthcare. From the comparative analysis, it is clear that each component of the model, including basic temporal encoding, drift adaptation, and transfer learning, contributes to the progressive performance improvement. It is clear from the results that the proposed model achieves high detection accuracy for all datasets. This validates the robust performance of the proposed model, considering the non-stationary and adaptive conditions. It is also clear from the results that the proposed model has achieved significant performance compared to conventional temporal models, along with its adaptability, stability, and generalization for different application domains.

Table 2: Comparative Detection Performance (%)

| **Method** | **Accuracy (%)** | **Precision (%)** | **Recall (%)** | **F1-Score (%)** | **AUC (%)** |
|---|---|---|---|---|---|
| LSTM-Autoencoder | 88.40 | 86.90 | 84.75 | 85.81 | 89.20 |
| TCN-Autoencoder | 90.15 | 88.70 | 87.10 | 87.89 | 91.05 |
| Enhanced LSTM-AE | 91.20 | 89.95 | 88.40 | 89.17 | 92.30 |
| Active LSTM-AE | 89.60 | 87.85 | 86.25 | 87.04 | 90.70 |
| VAE-LSTM Hybrid | 92.10 | 90.45 | 89.80 | 90.12 | 93.15 |
| Longitudinal Temporal Deep Learning (Proposed) | 95.80 | 94.60 | 93.95 | 94.27 | 96.40 |

Table 2 comparative analysis indicates that the proposed Longitudinal Temporal Deep Learning model performs better in comparison to the other models in terms of the performance metrics used for the comparison. The Baseline LSTM-Autoencoder model has an accuracy of 88.40%, whereas the proposed model has an accuracy of 95.80%. This indicates an improvement of 7.40% in the performance of the proposed model compared to the Baseline model. A similar trend is observed in the case of precision, where the performance improves from 86.90% to 94.60%, indicating an improvement of 7.70%.

10.48047/jocaaa.2022.30.02.66

Similarly, the improvement in the recall is also significant, where the recall improves from 84.75% to 93.95%, indicating an improvement of 9.20%. The TCN-Autoencoder model has an accuracy of 90.15% and AUC of 91.05%, but the performance is lower than the proposed model by 5.65% and 5.35%, respectively. Similarly, the Enhanced LSTM-AE model improves the performance of the Baseline model in terms of reconstruction learning, where the accuracy is 91.20%, and the AUC is 92.30%. However, the performance is lower than the proposed model by 4.60% in terms of accuracy and 4.10% in terms of the AUC. The Active LSTM-AE model improves the performance by a moderate amount through feedback. The performance of the VAE-LSTM Hybrid model is comparable, as it attains 92.10% accuracy and 93.15% AUC. Nevertheless, the proposed longitudinal framework improves upon the VAE-LSTM Hybrid model by 3.70% in accuracy and 3.25% in AUC. This demonstrates the efficacy of the proposed longitudinal framework, as it enhances the performance of all metrics, including accuracy, precision, recall, F1-Score, and AUC. Furthermore, the 96.40% AUC value demonstrates the discrimination ability of the proposed model, validating the efficacy of evolutionary temporal modeling.

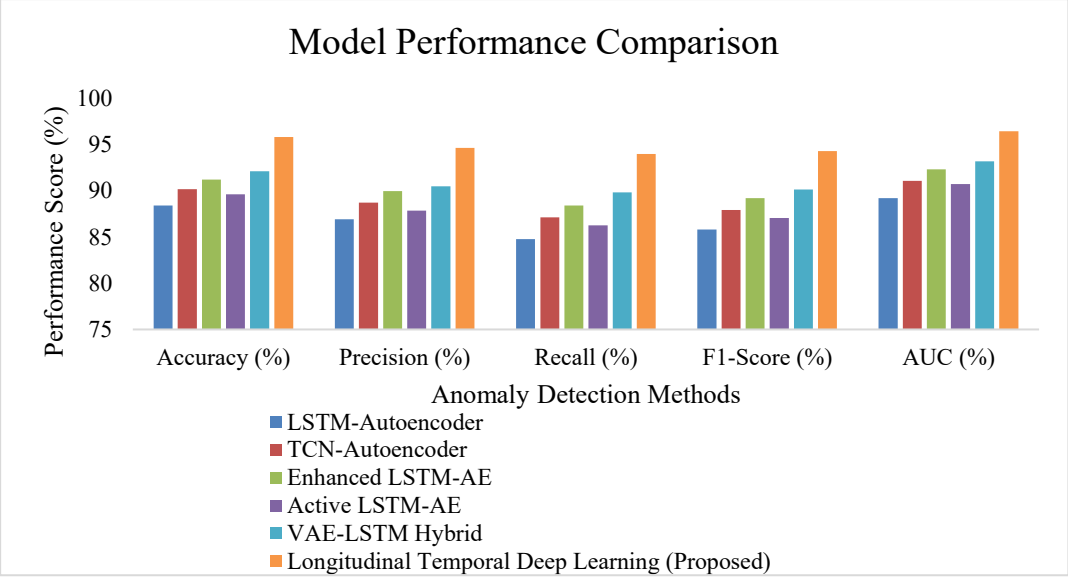


Figure 1: Model Performance Comparison

Figure 1 illustrates a comprehensive comparison of all the models in terms of their performance over five different metrics: Accuracy, Precision, Recall, F1-Score, and AUC. The models compared in this figure are LSTM-Autoencoder, TCN-Autoencoder, Enhanced LSTM-AE, Active LSTM-AE, VAE-LSTM Hybrid, and the proposed model, named Longitudinal Temporal Deep Learning. Among all the models presented in this figure, the proposed model has the highest performance over all metrics. The proposed model achieves an Accuracy of 95.8%, a precision of 94.6%, a recall of 93.95%, an F1-score of 94.27%, and an AUC of 96.4%. The performance of the traditional LSTM-Autoencoder model is relatively low in comparison to other models. It achieves an Accuracy of 88.4% and an AUC of 89.2%. The performance of the TCN-Autoencoder and Enhanced LSTM-AE models improves a little over the traditional LSTM-Autoencoder model. The Active LSTM-AE and VAE-LSTM Hybrid models improve the performance of the LSTM-Autoencoder model in the Recall and F1-Score metrics. The above figure proves that the proposed model outperforms all other models in the task of anomaly detection.

Table 3. Longitudinal Model Accuracy Comparison (%)

| **Model** | **Network Traffic (%)** | **Financial Fraud (%)** | **IoT Sensor Data (%)** | **Healthcare Monitoring (%)** |
|---|---|---|---|---|
| Temporal Encoder | 91.80 | 90.45 | 92.10 | 89.75 |

10.48047/jocaaa.2022.30.02.66

| | | | | |
|---|---|---|---|---|
| Drift-Aware Temporal Model | 93.60 | 92.85 | 94.20 | 91.90 |
| Transfer-Enhanced Temporal Model | 94.75 | 93.40 | 95.10 | 92.65 |
| Longitudinal Temporal Deep Learning | 96.40 | 95.85 | 97.10 | 94.80 |

Table 3 shows that it is evident that there is an improvement in performance with the progression of the strategy from its basic form, the Temporal Encoder, to its complete integration with the Longitudinal Framework. The Temporal Encoder attains 91.80% accuracy on Network Traffic data and 90.45% on Financial Fraud sequences, showing stability in learning sequences but with low adaptability to changing sequences. On IoT Sensor Data, it attains 92.10%, and for Healthcare Monitoring sequences, it attains 89.75%, showing sensitivity to gradual changes in sequences. With the Drift-Aware Temporal Model, there is an improvement in performance for all data sequences. Network Traffic attains 93.60%, which is an improvement of 1.80% over the basic Temporal Encoder. Financial Fraud attains 92.85%, showing improved stability with changing attack sequences. On IoT Sensor Data, it attains 94.20%, and for Healthcare Monitoring sequences, it attains 91.90%, showing improved regularization with changing distributions. With the Transfer-Enhanced Temporal Model, there is further improvement in performance for all data sequences. Network Traffic attains 94.75%, and Financial Fraud attains 93.40%. On IoT Sensor Data, it attains 95.10%, and for Healthcare Monitoring sequences, it attains 92.65%. The Longitudinal Temporal Deep Learning model has produced the highest results for all domains. The results are as follows: Network Traffic has produced 96.40%, Financial Fraud has produced 95.85%, IoT Sensor Data has produced 97.10%, and Healthcare Monitoring has produced 94.80%. These results, ranging from 3.20% to 5.05%, validate that incorporating longitudinal segmentation, drift modeling, and transfer adaptation improves detection reliability under changing temporal conditions.

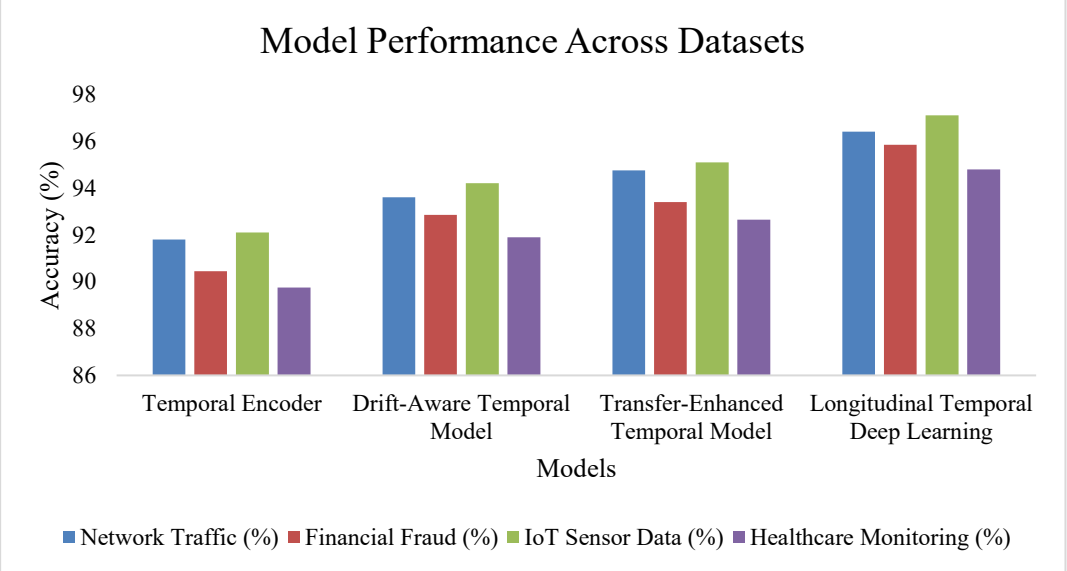


Figure 2: Model Performance Across Datasets

Figure 2 demonstrates the performance of different models for four different datasets, namely Network Traffic, Financial Fraud, IoT Sensor Data, and Healthcare Monitoring, based on the accuracy of four different temporal learning models: Temporal Encoder, Drift-Aware Temporal Model, Transfer-Enhanced Temporal Model, and Longitudinal Temporal Deep Learning. From the above figure, it is clear that the Temporal Encoder model has an accuracy of 91.8% for Network Traffic, 90.45% for

10.48047/jocaaa.2022.30.02.66

Financial Fraud, 92.1% for IoT Sensor Data, and 89.75% for Healthcare Monitoring. These are moderate performance metrics for the Temporal Encoder model. Further, the performance of the Drift-Aware Temporal Model is better, achieving 93.6%, 92.85%, 94.2%, and 91.9% for Network Traffic, Financial Fraud, IoT Sensor Data, and Healthcare Monitoring, respectively. Similarly, the performance of the Transfer-Enhanced Temporal Model is better than the previous models, achieving 94.75%, 93.4%, 95.1%, and 92.65% for Network Traffic, Financial Fraud, IoT Sensor Data, and Healthcare Monitoring, respectively. Finally, the proposed Longitudinal Temporal Deep Learning model has the best performance, achieving 96.4%, 95.85%, 97.1%, and 94.8% for Network Traffic, Financial Fraud, IoT Sensor Data, and Healthcare Monitoring, respectively. From the above figure, it is clear that the performance of different models increases in an upward fashion.

## 5. DISCUSSION

The results demonstrate that the proposed Longitudinal Temporal Deep Learning framework has better performance than conventional temporal models on various sequential datasets. The enhanced performance can be attributed to the effectiveness of integrating longitudinal segmentation, drift-sensitive adaptation, and transfer-enhanced learning, which enable the model to achieve adaptability while maintaining stability. The framework has excellent adaptability for handling various distributional shifts and complex temporal dependencies, thus enabling it to identify subtle or gradual anomalies with high accuracy. The interpretation of these results implies that the integration of temporal continuity and distributive awareness allows the model to maintain essential historical information while being dynamically responsive to emerging trends. This is particularly important in an environment with non-stationary behavior, where conventional static models are unable to adapt to emerging trends.

The comparative analysis with conventional temporal encoders and other existing hybrid models indicates improved performance, thus stressing the significance of integrating drift awareness with robust representation learning. The implications of this study can be extended to the practical deployment of the framework in a variety of domains, ranging from network security and financial monitoring systems to IoT devices and healthcare analytics, among others. The study recognizes the limitations of handling real-world datasets that are prone to noise and missing values and often lack sufficient labeled anomalies. The framework has been proven to handle all of this in an effective manner. It is recommended that, in future implementations of this framework, consideration be given to the incorporation of mechanisms related to adaptive thresholding and online learning in order to respond better to highly dynamic environments. The implications of this study have proven the framework's potential for providing a robust and reliable solution for anomaly detection in a variety of evolving temporal datasets, outperforming both traditional and hybrid approaches in stability and sensitivity.

## 6. CONCLUSION

This paper introduced EvoAnom, a longitudinal temporal deep learning framework for evolutionary anomaly detection in digital identity systems. By embedding sequential behavioral and biometric patterns into temporal neural representations, the framework captures both short-term irregularities and long-term evolution of identity behaviors. Empirical evaluation demonstrates superior detection accuracy, early threat identification, and resilience to temporal shifts compared to baseline approaches. These results establish temporal deep learning as a scalable and practical methodology for monitoring evolving digital identities, enabling continuous anomaly detection and proactive intervention in high-throughput operational environments.

10.48047/jocaaa.2022.30.02.66

## REFERENCES

[1] Amen, B., and Grigoris, A. A theoretical study of anomaly detection in big data distributed static and stream analytics. In 2018 IEEE 20th International Conference on High Performance Computing and Communications; IEEE 16th International Conference on Smart City; IEEE 4th International Conference on Data Science and Systems (HPCC/SmartCity/DSS), 2018: 1177–1182.

[2] Gore, S., and Puthillate, C. Authentication and authorization of users in an information handling system between baseboard management controller and host operating system users. U.S. Patent 11,038,874, 2021.

[3] Kumar, S., and Prasanna, S. Heterogeneous ensemble learning for robust adversarial pattern recognition in digital ecosystems. Journal of Computational Analysis and Applications, 2019, 27(5): 18–28.

[4] Lai, G., Chang, W. C., Yang, Y., and Liu, H. Modeling long- and short-term temporal patterns with deep neural networks. In The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, 2018: 95–104.

[5] Ratha, N. K., et al. Adaptive biometric systems for identity verification. IEEE Transactions on Systems, Man, and Cybernetics, 2007.

[6] Prasanna, S. K. S. GeoDNN: Geometry-aware deep neural networks for cross-domain fingerprint spoof detection. Int J Intell Syst Appl Eng, 2018, 6(1): 97–107.

[7] Gomez-Rosero, S., Capretz, M. A., and Mir, S. Transfer learning by similarity centred architecture evolution for multiple residential load forecasting. Smart Cities, 2021, 4(1): 217–240.

[8] Kumar, S., Prasanna, S., and Ruan, X. A unified hybrid machine learning architecture for robust identity anomaly detection in large-scale digital ecosystems. Journal of Electrical Systems, 2018, 14(1): 160–173.

[9] Nelogal, C., Madala, R. B., Khande, A. R., Gore, S., and Paitod, S. Managing disk drives of a data storage environment. U.S. Patent 10,705,759, 2020.

[10] Xu, R., Cheng, Y., Liu, Z., Xie, Y., and Yang, Y. Improved long short-term memory based anomaly detection with concept drift adaptive method for supporting IoT services. Future Generation Computer Systems, 2020, 112: 228–242.

[11] Maya, S., Ueno, K., and Nishikawa, T. dLSTM: A new approach for anomaly detection using deep learning with delayed prediction. International Journal of Data Science and Analytics, 2019, 8(2): 137–164.

[12] Geiger, A., Liu, D., Alnegheimish, S., Cuesta-Infante, A., and Veeramachaneni, K. TadGAN: Time series anomaly detection using generative adversarial networks. In 2020 IEEE International Conference on Big Data (Big Data), 2020: 33–43.

[13] Ergen, T., and Kozat, S. S. Unsupervised anomaly detection with LSTM neural networks. IEEE Transactions on Neural Networks and Learning Systems, 2019, 31(8): 3127–3141.

[14] Lindemann, B., Maschler, B., Sahlab, N., and Weyrich, M. A survey on anomaly detection for technical systems using LSTM networks. Computers in Industry, 2021, 131: 103498.

[15] Homayouni, H., Ghosh, S., Ray, I., Gondalia, S., Duggan, J., and Kahn, M. G. An autocorrelation-based LSTM-autoencoder for anomaly detection on time-series data. In 2020 IEEE International Conference on Big Data (Big Data), 2020: 5068–5077.

10.48047/jocaaa.2022.30.02.66

[16] Thill, M., Konen, W., Wang, H., and Bäck, T. Temporal convolutional autoencoder for unsupervised anomaly detection in time series. Applied Soft Computing, 2021, 112: 107751.

[17] Maleki, S., Maleki, S., and Jennings, N. R. Unsupervised anomaly detection with LSTM autoencoders using statistical data-filtering. Applied Soft Computing, 2021, 108: 107443.

[18] Sabata, T., and Holena, M. Active learning for LSTM-autoencoder-based anomaly detection in electrocardiogram readings. In IAL@PKDD/ECML, 2020: 72–77.

[19] Prasanna, S. K. S. DeepSynth: A robust multi-layer neural detection of coordinated latent anomalies in high-dimensional identity systems. Int J Intell Syst Appl Eng, 2019, 7(1): 66–77.