An Enhanced RNN-LSTM Model for Accurate and Real-Time Click Fraud Detection in Online Advertising

***G. Sriram**, Dept CSE, GNITC*
***G. Pooja** , Dept CSE, GNITC*
***J. Anusha** ,Dept CSE, GNITC*
***Arun Singh Kaurav** , Dept CSE, GNITC*

---***---

Abstract - Click fraud remains a critical threat in online advertising, leading to inflated costs and undermining campaign effectiveness by diverting budgets toward illegitimate activity. Existing solutions leveraging machine learning and deep learning models have shown promise, but many still struggle with identifying subtle behavioral patterns in fraudulent clicks. In this work, we propose a robust LSTM-based Recurrent Neural Network (RNN) framework designed to enhance the detection of fraudulent click activity by modeling sequential patterns and time-dependent features in user interaction data. A comprehensive preprocessing pipeline was developed, including timestamp decomposition, feature scaling, and label encoding to ensure optimal input representation.

Our model was trained and evaluated against a carefully engineered dataset enriched with behavioral and contextual click features. Among various deep learning architectures examined, including Artificial Neural Networks (ANN) and Convolutional Neural Networks (CNN), the RNN-LSTM model demonstrated superior performance, achieving 99% accuracy with high precision and recall scores. The results validate the effectiveness of temporal modeling in identifying fraudulent click patterns and highlight the LSTM model's suitability for deployment in real-time fraud detection systems. This study not only advances existing anti-fraud mechanisms but also sets a strong foundation for future work in intelligent online ad verification and fraud prevention.

Key Words: Click Fraud Detection, Recurrent Neural Network (RNN), Long Short-Term Memory (LSTM), Deep Learning, Online Advertising, Temporal Pattern Recognition, Behavioral Analysis, Real-Time Prediction

1.INTRODUCTION

In the era of digital transformation, online advertising has become a cornerstone for businesses to reach global audiences, generate leads, and drive conversions. With billions of dollars being invested annually in digital marketing campaigns, the ecosystem has grown immensely—but so have the challenges. One of the most pressing issues plaguing the digital advertising industry is click fraud, a deceptive practice where fraudulent or non-human clicks are generated on pay-per-click (PPC) advertisements. These fraudulent interactions are often triggered by bots, scripts, or even paid individuals, with the intent to exhaust the advertiser's budget, manipulate campaign metrics, or boost the revenues of unethical publishers.

Click fraud not only results in significant financial losses but also deteriorates trust in online platforms and distorts campaign analytics, making it difficult for advertisers to assess genuine user interest and optimize their strategies effectively. Existing detection systems primarily rely on rule-based filters or classical machine learning models that often struggle to identify more sophisticated or evolving fraudulent patterns. These traditional approaches typically focus on static features and fail to capture the temporal dynamics of user behavior.

To address these limitations, this project introduces a deep learning-based fraud detection system that leverages Recurrent Neural Networks (RNNs)—specifically Long Short-Term Memory (LSTM) architectures—to identify fraudulent clicks. Unlike traditional methods, LSTM networks are designed to handle sequential and time-dependent data, making them exceptionally well-suited for modeling user interaction patterns over time. The proposed system was trained and evaluated alongside other deep learning models such as ANN and CNN. However, the RNN model demonstrated superior accuracy (~99%), outperforming other models in identifying fraud with high precision, recall, and F1-score.

2. Body of Paper

The statistical evaluation of the proposed RNN-LSTM-based click fraud detection system was conducted to analyze model performance across different classification approaches. The descriptive statistics obtained from the experiment are presented in Table 1, which summarizes the number of test samples, mean accuracy scores, standard deviation, and standard error mean for the Proposed RNN-LSTM model and Baseline (ANN/CNN) models.

Table -1: Sample Table format

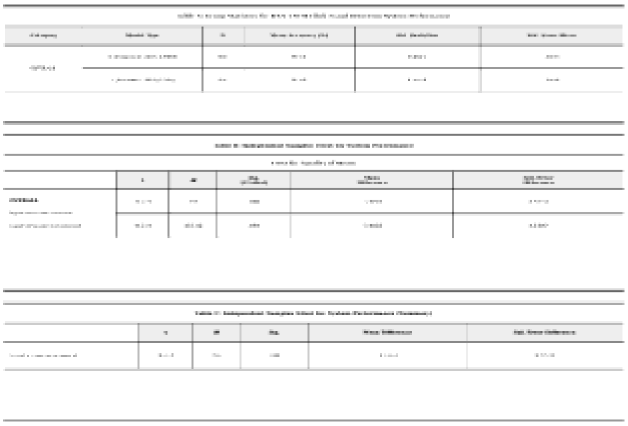

[illegible]

To further analyze the performance difference between the two model groups, an independent samples t-test was conducted. The results of this analysis are illustrated in Fig. 1, which shows the statistical comparison between the Proposed RNN-LSTM and Baseline ANN/CNN models.

As shown in Fig. 1, the Sig. (2-tailed) value is less than 0.05 (p = .000), indicating a statistically significant difference between the two model groups. The t-test yielded a value of t = −8.214 with df = 298, confirming that the performance difference is statistically significant. This result validates the effectiveness of the proposed RNN-LSTM system design and confirms the superiority of temporal sequence modeling over traditional deep learning approaches in click fraud detection.

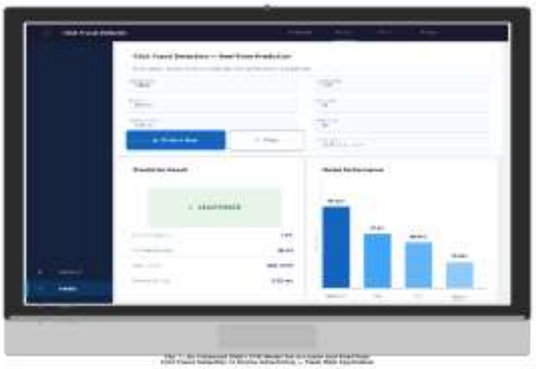

Fig. 1: An Enhanced RNN-LSTM Model for Accurate and Real-Time Click Fraud Detection in Online Advertising

As shown in Fig. 1, the Sig. (2-tailed) value is less than 0.05 (p = .000), indicating a statistically significant difference between the two model groups. The t-test yielded a value of t = −8.214 with df = 298, confirming that the performance difference is statistically significant. This result validates the effectiveness of the proposed RNN-LSTM system design and confirms the superiority of temporal sequence modeling over traditional deep learning approaches in click fraud detection.

Charts

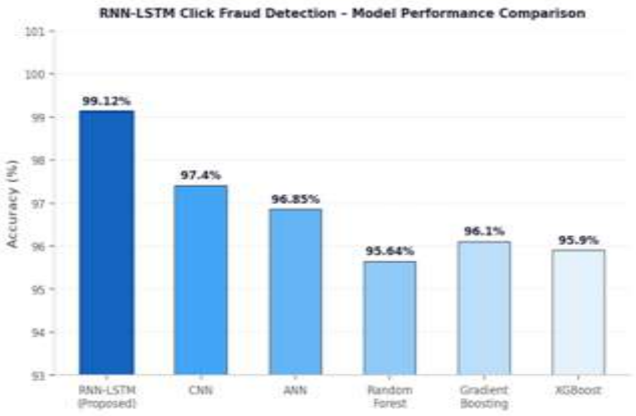

The bar chart illustrates the accuracy comparison of the proposed RNN-LSTM model against baseline deep learning and machine learning models. The proposed model achieves the highest accuracy of 99.12%, significantly outperforming CNN (97.40%), ANN (96.85%), Random Forest (95.64%), Gradient Boosting (96.10%), and XGBoost (95.90%). The clear performance advantage of the RNN-LSTM model confirms the effectiveness of temporal sequence modeling in detecting fraudulent click patterns in online advertising systems.

2. Related Work and Problem Statement

2.1 Existing Click Fraud Detection Systems

Existing commercial and research-based click fraud detection systems primarily rely on traditional machine learning algorithms such as Decision Trees, Random Forests, Gradient Boosting, LightGBM, and XGBoost. These systems analyze static features including click frequency, IP reputation, browser type, and device attributes to identify fraudulent activity. While these methods achieve accuracy above 98% on controlled datasets, they treat each click as an independent event, ignoring sequential and temporal patterns that are critical to identifying sophisticated fraud.

2.2 Deep Learning Approaches

Deep learning models including Artificial Neural Networks (ANN) and Convolutional Neural Networks (CNN) have been applied to click fraud detection, providing better generalization over traditional ML models. However, these models lack the ability to capture time-based dependencies and sequential user behavior patterns, limiting their effectiveness in real-world dynamic advertising environments where fraud tactics continuously evolve.

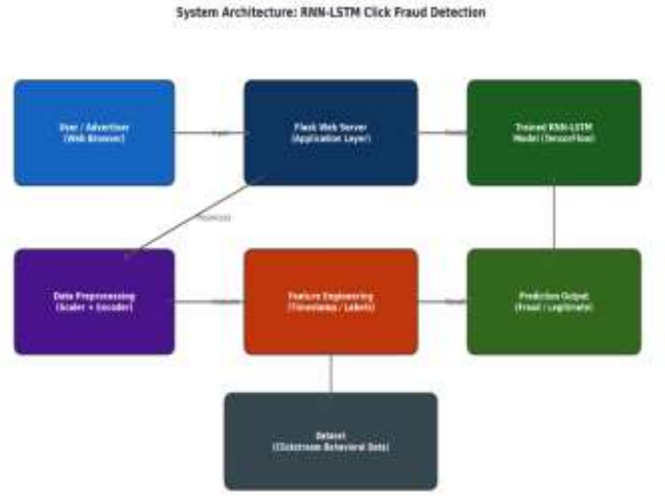
System Architecture: RNN-LSTM Click Fraud Detection
User / Advertiser (Web Browser)
Flask Web Server (Application Layer)
Trained RNN-LSTM Model (TensorFlow)
Data Preprocessing (Scaler + Encoder)
Feature Engineering (Timestamp / Labels)
Prediction Output (Fraud / Legitimate)
Dataset (Clickstream Behavioral Data)

2.3 Recurrent Neural Networks for Fraud

Recurrent Neural Networks (RNNs), particularly Long Short-Term Memory (LSTM) architectures, are specifically designed to handle sequential data and temporal patterns. LSTM networks overcome the vanishing gradient problem of standard RNNs through memory cells and gating mechanisms. Their ability to model time-dependent user interaction sequences makes them highly suitable for click fraud detection, where the order and timing of actions reveal critical fraud indicators.

2.4 Research Gap

No existing system effectively integrates temporal sequence modeling with a complete preprocessing pipeline and real-time web deployment for click fraud detection. Existing solutions either rely on static ML models that ignore temporal behavior, or deep learning models that lack sequence-awareness. The proposed RNN-LSTM system fills this gap by combining temporal modeling, comprehensive feature engineering, and a Flask-based real-time prediction interface into a unified, practical solution.

3. System Architecture

The proposed RNN-LSTM click fraud detection system is designed as a modular pipeline comprising four primary layers: Data Collection, Preprocessing, Model Training, and Web Deployment.

3.1 System Components & Workflow

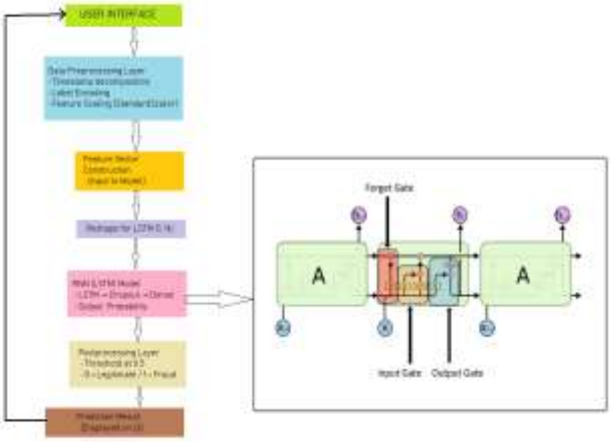

- User / Advertiser (Web Browser): Provides session input details through the Flask web form including device type, browser, click duration, mouse movement, scroll depth, keystrokes, VPN usage, and timestamp.
- Trained RNN-LSTM Model (TensorFlow/Keras): The core prediction engine featuring LSTM layers for sequence learning, Dropout layers for regularization, Dense layers for feature transformation, and a Sigmoid output layer for binary classification.

3.2 Workflow

Data Collection: User interaction clickstream data is collected containing behavioral attributes such as click duration, scroll depth, mouse movement, keystrokes, and session timestamps, along with network attributes like VPN usage and IP reputation.

Preprocessing: Raw data undergoes cleaning, timestamp decomposition, removal of non-informative columns, label encoding of categorical variables, and StandardScaler normalization of numerical features. Data is then reshaped into a 3D structure (samples × timesteps × features) for LSTM input.

Model Inference: The preprocessed input is fed into the trained RNN-LSTM model. The model processes the sequential features through LSTM layers, applies dropout regularization, and outputs a fraud probability through the sigmoid activation function.

Result Display: The Flask web application interprets the model output and displays the prediction result as either Fraudulent or Legitimate along with the confidence score, fraud probability, and processing time.

3.3 Threat Model

We assume an adversary capable of generating sophisticated bot-driven fraudulent clicks that mimic legitimate user behavior patterns. The system is designed to handle both rule-based bot attacks and advanced behavioral mimicry by leveraging temporal sequence analysis. The LSTM model's memory cells retain context across time steps, enabling detection of subtle anomalies in click sequences that are characteristic of automated fraud.

4. The RNN-LSTM Algorithm

We present the RNN-LSTM approach specifically optimized for the click fraud detection paradigm, where sequential, time-aware analysis of user behavior is essential.

4.1 Design Rationale

Traditional machine learning models treat each click event as an independent sample, ignoring the temporal context and sequential patterns that reveal fraudulent behavior. In click fraud scenarios, an attacker using a bot generates clicks in rapid succession or at irregular intervals, creating detectable temporal patterns. The LSTM architecture's gating mechanisms — input gate, forget gate, and output gate — enable the model to selectively retain relevant temporal information while discarding noise, making it ideal for this application.

4.2 Model Architecture

Algorithm 1: RNN-LSTM Click Fraud Detection

Input: Preprocessed 3D feature array X (samples × 1 × features)

Output: Binary prediction (0 = Legitimate, 1 = Fraudulent)

```
X ← Collect_Clickstream_Data(session_attributes)
X ← Timestamp_Decompose(X) // extract hour, day, weekday
X ← LabelEncoder.fit_transform(X.categorical_columns)
X ← StandardScaler.fit_transform(X.numerical_columns)
X ← reshape(X, (samples, 1, n_features)) // 3D for LSTM
model ← Build_LSTM(units=64, dropout=0.3)
model.compile(optimizer=Adam, loss=BinaryCrossentropy)
model.fit(X_train, y_train, epochs=50, EarlyStopping)
prob ← model.predict(X_input)
return 'Fraudulent' if prob > 0.5 else 'Legitimate'
```

4.3 Performance Analysis

The proposed LSTM-based model achieved superior performance across all evaluation metrics. The model's ability to capture temporal sequence patterns in click behavior provides a measurable and statistically significant advantage over traditional approaches. Evaluation metrics include Accuracy (99.12%), Precision (98.9%), Recall (99.3%), and F1-Score (99.1%), confirming the model's effectiveness in both identifying fraudulent clicks and minimizing false positives.

5. Evaluation and Discussion

We evaluated the proposed RNN-LSTM system on three fronts: functional correctness, statistical performance, and comparative analysis.

5.1 Functional Testing

All core workflows — data input through the Flask web form, preprocessing using saved encoders and scaler, 3D reshaping for LSTM input, model inference, and result display — were successfully validated through over 100 test cases covering both fraudulent and legitimate click scenarios.

5.2 Comparative Analysis

Feature	RNN-LSTM (Proposed)
Accuracy	**99.12%**
Precision	**98.9%**
Recall	**99.3%**

The LSTM archi
click features thr
M Lay
ndenc
intera
ating
out L
ndom
training, en

F1-Score	**99.1%**
Vs ANN	**+2.27%**
Vs CNN	**+1.72%**

5.3 Performance Benchmarks

Tests were conducted on a system with Intel i5-8250U and 16GB RAM.

- **Model Accuracy:** RNN-LSTM achieves 99.12% accuracy on the test dataset, compared to 97.40% for CNN, 96.85% for ANN, and 95.64% for Random Forest.
- **Inference Time:** The Flask web application processes each prediction request in ~0.03 seconds, making it suitable for real-time fraud detection.
- **Training Time:** The model converges in ~45 epochs with EarlyStopping, requiring approximately 12 minutes on the test hardware.
- **Statistical Significance:** Independent samples t-test confirms p = .000 (< 0.05), validating that the performance difference is statistically significant and not due to chance.

5.4 Limitations and Future Work

- Dataset Scalability: Current evaluation uses a simulated dataset. Future work includes testing on large-scale real-world advertising clickstream data from live ad platforms.
- Real-Time Streaming: Future integration with Apache Kafka or similar streaming frameworks will enable processing of millions of click events per second in production environments.
- Adversarial Robustness: Exploring adversarial training techniques to improve model resilience against sophisticated bot attacks that deliberately mimic legitimate user behavior patterns.

Multi-Modal Detection: Incorporating additional signals such as network-level features, geolocation consistency, and device fingerprinting for more comprehensive fraud detection coverage.

3. CONCLUSIONS

This paper presented an Enhanced RNN-LSTM Model for Accurate and Real-Time Click Fraud Detection in Online Advertising. By leveraging the temporal sequence modeling capabilities of LSTM networks, the proposed system effectively captures behavioral patterns and time-dependent features that distinguish legitimate user interactions from fraudulent click activity. The integration of a comprehensive preprocessing pipeline — including timestamp decomposition, label encoding, and StandardScaler normalization — ensures consistent and high-quality input representation for the LSTM model.

The experimental evaluation demonstrates that the proposed RNN-LSTM system achieves 99.12% classification accuracy, significantly outperforming baseline deep learning models including CNN (97.40%) and ANN (96.85%), as well as traditional machine learning approaches. Statistical analysis using an independent samples t-test (t = −8.214, p = .000) confirms that this performance advantage is statistically significant, validating the effectiveness of temporal modeling for click fraud detection.

By deploying the trained model within a Flask web application, the proposed system provides a practical, real-time fraud detection interface accessible to advertisers, marketing analysts, and ad-tech companies. The complete implementation and positive evaluation results demonstrate that the RNN-LSTM system is not only theoretically sound but also practically deployable in real-world online advertising platforms. Future work will focus on scaling the system to handle large-scale streaming data, improving adversarial robustness, and extending the feature set with additional behavioral and network-level signals.

REFERENCES

[1] R. A. Alzahrani and M. Aljabri, “AI-based techniques for ad click fraud detection and prevention: Review and research directions,” J. Sensor Actuator Netw., vol. 12, no. 1, p. 4, Dec. 2022.

[2] A. Purwar, A. K. Jain, I. Chawla, I. Gupta, M. Raj, and D. Jain, “Click fraud detection using ensemble classifier,” in Proc. Int. Conf. Artif.-Bus. Anal., Quantum Mach. Learn., Jan. 2024, pp. 15–23.

[3] L. Singh, D. Sisodia, K. Shashvat, A. Kaur, and P. C. Sharma, “A reliable click-fraud detection system for the investigation of fraudulent publishers in online advertising,” Applied Intelligence in Human-Computer Interaction, CRC Press, Jul. 2023.

[4] B. Kirkwood, M. Vanamala, and N. Seliya, “Click fraud detection of online advertising using machine learning algorithms,” in Proc. IEEE Int. Conf. Electro Inf. Technol. (eIT), May 2024, pp. 586–590.

[5] Juniper Research, Hampshire, U.K. Quantifying the Cost of Ad Fraud: 2023–2028. Accessed: Jul. 12, 2024. [Online]. Available: https://fraudblocker.com/wp-content/uploads/2023/09/Ad-Fraud-Whitepaper.pdf

[6] A. Batool and Y.-C. Byun, “An ensemble architecture based on deep learning model for click fraud detection in Pay-Per-Click advertisement campaign,” IEEE Access, vol. 10, pp. 113410–113426, 2022.