iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** VIII **Month of publication:** August 2025

DOI: https://doi.org/10.22214/ijraset.2025.73785

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue VIII Aug 2025- Available at www.ijraset.com

An Ensemble Architecture Based on Deep Learning Model for Click Fraud Detection in Pay-Per-Click Advertisement Campaign

Ch.Divya[1], Nalamalapu Gayathri[2], Gutti Bhagyalakshmi[3], Devathoti Pavithra[4], Metla Devi Aditya Prasad[5], Nallapaneni Manil[6]

[1]*Assistant Professor,* [2,3,4,5,6]*UnderGraduate, CSE-Data Science Department, St. Ann's College of Engineering & Technology, Chirala, Andhra Pradesh*

Abstract: ***Pay-per-click (PPC) advertising has become a foundational component of digital marketing, enabling advertisers to reach targeted audiences through keyword-based ad placements. However, the rise in automated scripts, bots, and click farms has led to a surge in click fraud—invalid clicks that artificially inflate advertising costs and mislead performance metrics. Existing fraud detection techniques often rely on static rule-based systems or shallow machine learning models, which are inadequate in identifying evolving and obfuscated fraudulent behavior. This paper proposes an ensemble deep learning architecture that integrates a Convolutional Neural Network (CNN) with a Long Short-Term Memory (LSTM) network for robust click fraud detection. The CNN module captures local temporal patterns within clickstream features, while the LSTM module models long-term behavioral dependencies across user sessions. Experimental results on benchmark datasets demonstrate that the proposed hybrid model outperforms conventional models in accuracy, precision, recall, and F1-score. The ensemble approach effectively reduces false positives while adapting to evolving fraud signatures, offering a scalable and intelligent solution for securing PPC platforms.***

Keywords: Click fraud, PPC advertising, deep learning, ensemble model, CNN-LSTM, fraud detection, ad security.

I. INTRODUCTION

The proliferation of digital advertising has reshaped the marketing landscape, with Pay-Per-Click (PPC) campaigns becoming a central strategy for online outreach and revenue generation. PPC models allow advertisers to bid on relevant keywords, paying only when users click on their ads. While this model provides cost-efficiency and performance measurability, it is also highly susceptible to fraudulent exploitation in the form of click fraud—a deceptive practice involving artificial inflation of ad clicks to drain advertiser budgets or distort competitor performance.

Click fraud can be executed through various means, including botnets, click farms, and automated scripts. These fraudulent clicks are often indistinguishable from legitimate user behavior, especially when attackers attempt to mimic human patterns. Traditional click fraud detection systems rely heavily on rule-based mechanisms or shallow machine learning classifiers that analyze features like IP addresses, session durations, and click intervals.

However, these methods lack adaptability, struggle with concept drift, and often yield high false positive rates when exposed to evolving adversarial behaviors.

With the advent of deep learning, more sophisticated detection models have been developed to learn complex, non-linear representations from high-dimensional data. Convolutional Neural Networks (CNNs) excel at extracting spatial features and have been applied to sequence-like user behavior data. Meanwhile, Long Short-Term Memory (LSTM) networks are effective at modeling temporal dependencies and have shown success in fraud and anomaly detection tasks. Despite these advances, single-model approaches often underperform in real-world conditions due to their limited perspective on feature hierarchies or temporal consistency.

This study proposes a robust ensemble architecture combining CNN and LSTM models for enhanced click fraud detection in PPC systems. The CNN module captures local behavior patterns across short time windows, while the LSTM component processes long-term dependencies across user interaction sessions. By fusing these complementary perspectives, the proposed model achieves high accuracy and robustness, even in the presence of noisy or obfuscated data.

©IJRASET: All Rights are Reserved

The key contributions of this paper are as follows:

- Development of a hybrid deep learning model that combines CNN and LSTM layers for adaptive click fraud detection.
- Feature extraction and preprocessing techniques tailored for temporal and session-based clickstream data.
- Comparative performance evaluation against baseline models, demonstrating improved detection rates and reduced false alarms.
- A modular implementation framework suitable for integration into real-time PPC fraud monitoring systems.

By leveraging the strengths of deep learning architectures in both feature representation and temporal modeling, the proposed ensemble model provides a scalable, intelligent solution to the growing challenge of click fraud in digital advertising ecosystems.

II. LITERATURE SURVEY

Click fraud has emerged as a significant threat to the integrity and profitability of online advertising ecosystems. Over the past decade, extensive research has been conducted on fraud detection strategies, ranging from rule-based techniques to advanced machine learning and deep learning models. This section reviews prior works that have contributed to the detection of click fraud in PPC environments, highlighting the evolution of methodologies and the growing relevance of ensemble deep learning architectures.

A. Traditional Detection Approaches

Early click fraud detection methods were primarily rule-based systems that used fixed thresholds on IP addresses, click timestamps, user agents, and geographic location. These methods flagged users based on anomalies such as multiple clicks from the same IP or unusually short dwell times. Although computationally efficient, rule-based systems lacked adaptability and were vulnerable to evasion through IP masking and user behavior simulation.

Statistical models such as regression analysis and Naïve Bayes classifiers were later introduced to model click distributions and predict suspicious activity. However, their performance was hindered by the linearity assumption and inability to capture non-obvious fraud patterns.

B. Machine Learning Models

Supervised machine learning models, including Decision Trees, Support Vector Machines (SVM), Random Forests, and Gradient Boosting, improved detection performance by learning from labeled data. These algorithms could capture complex feature interactions and achieved moderate success in classification tasks. Nevertheless, they were limited by their reliance on feature engineering and struggled with concept drift, where fraud patterns evolve over time.

Unsupervised techniques such as k-means clustering and Isolation Forests were also explored for anomaly detection in unlabeled clickstream datasets. While useful in identifying novel fraud types, these models lacked interpretability and often generated high false-positive rates.

C. Deep Learning Techniques in Fraud Detection

Recent studies have explored deep learning as a viable solution to overcome the shortcomings of traditional models. CNNs have been employed to extract hierarchical patterns from structured and sequential data, particularly in user interaction logs and session traces. LSTMs, a type of recurrent neural network (RNN), have shown effectiveness in learning long-range dependencies in time-series data such as click sequences, session transitions, and dwell time trajectories.

For instance, Liu et al. [1] implemented a CNN for fraud detection in mobile ad networks by converting clickstream data into spatial grids. Similarly, Hu et al. [2] utilized LSTM networks to model user behavioral sequences, capturing temporal fraud indicators such as repetitive clicking patterns and short-lived sessions.

Despite promising results, single-model architectures often struggle to generalize across diverse fraud scenarios and dataset imbalances. To overcome this, ensemble models that combine the strengths of multiple deep learning networks have been proposed.

D. Ensemble Deep Learning Models

Ensemble methods, particularly those combining CNN and LSTM layers, offer a hybrid perspective—where CNNs process localized features and LSTMs handle temporal dependencies. These architectures have been applied in domains such as credit card fraud, network intrusion detection, and medical diagnosis, demonstrating superior classification performance and robustness to noisy inputs.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

In the context of PPC fraud, ensemble deep models remain underexplored. Existing studies often treat spatial and temporal patterns in isolation, missing the potential synergy of multi-perspective learning. Furthermore, most models are not optimized for real-time inference or integration into live ad systems.

E. Research Gap

Despite advancements in deep learning and behavioral modeling, the following limitations persist in current click fraud detection literature:

- Overreliance on single-model architectures that lack generalization
- Poor adaptability to dynamic and adversarial click behaviors
- Inadequate integration of spatial and temporal feature representations
- Limited application of ensemble deep learning for PPC-specific fraud scenarios

The proposed ensemble CNN-LSTM model directly addresses these challenges by fusing spatial and temporal insights into a unified architecture, enabling adaptive and high-fidelity fraud detection in dynamic PPC environments.

III.PROPOSED METHODOLOGY

The proposed framework adopts a deep ensemble architecture that integrates a Convolutional Neural Network (CNN) and a Long Short-Term Memory (LSTM) network to detect click fraud in Pay-Per-Click (PPC) advertising systems. This hybrid architecture leverages the spatial learning capacity of CNN and the temporal sequence modeling capabilities of LSTM to accurately classify fraudulent versus legitimate clicks. The system pipeline comprises data preprocessing, feature transformation, model construction, training, and prediction.

A. Data Preprocessing

Raw clickstream data typically consists of user identifiers, timestamps, device details, session durations, IP addresses, and click behaviors.

The following preprocessing steps are applied to prepare the data for model ingestion:

- Missing Value Handling: Rows with missing critical fields such as timestamps or session IDs are discarded.
- Encoding Categorical Features: Attributes like browser type, device class, and location are converted using label encoding.
- Timestamp Conversion: Time-based features are converted into numerical form (e.g., hour of day, click interval).
- Normalization: Numerical features such as session time and click frequency are scaled using Min-Max normalization to fall within the [0, 1] range.
- Sequence Formation: Click events are grouped by session and sorted chronologically to form user activity sequences for LSTM input.

B. Feature Engineering

Key engineered features used in the model include:

- Number of clicks per session
- Average dwell time
- IP repetition rate
- Time since last click
- Device-user pair frequency
- Click-through rate deviation

These features are constructed to capture user behavior patterns that often signal fraudulent activity, such as rapid clicking, repeated IPs, and uniform session lengths.

C. CNN Module

The Convolutional Neural Network acts as a local feature extractor. It processes the structured input matrix formed from click features and identifies spatial dependencies across feature dimensions:

- Input: 2D matrix of shape (sequence_length, feature_dimension)

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

- Layers:
  - 1D Convolutional Layer with ReLU activation
  - Max Pooling Layer to reduce dimensionality
  - Batch Normalization for regularization

- Output: Feature map forwarded to the LSTM layer

The CNN enhances local pattern detection, such as burst clicks or repeated access from specific devices.

D. LSTM Module

The LSTM component captures the temporal evolution of click sequences. It models long-term dependencies such as user behavior across time windows or repeated fraudulent patterns:

- Input: Sequential output from CNN or raw time-sorted click sequences
- Layers:
  - LSTM Layer with 128 hidden units
  - Dropout layer to prevent overfitting

- Output: Final hidden state vector fed to the dense output layer

The LSTM enables dynamic analysis of how user actions unfold over time, improving the system's ability to detect complex and evolving fraud behaviors.

E. Output Layer and Classification

The final layer is a fully connected dense layer followed by a softmax or sigmoid activation depending on the classification type:

- Binary Classification: Sigmoid activation with threshold at 0.5
- Multiclass (if extended): Softmax for multiple fraud types

The model is trained using binary cross-entropy loss and optimized using the Adam optimizer with a learning rate of 0.001.

F. Model Training Strategy

- Train-Test Split: 80:20 ratio
- Batch Size: 128
- Epochs: 25–50 depending on convergence
- Regularization: Dropout (0.5), Early stopping based on validation loss

The model is trained using TensorFlow/Keras and evaluated using accuracy, precision, recall, F1-score, and AUC-ROC metrics.

IV. IMPLEMENTATION

The proposed CNN-LSTM ensemble model was implemented using Python with deep learning frameworks that support modular, scalable experimentation. The implementation involved dataset handling, feature engineering, model architecture construction, training procedures, and evaluation mechanisms.

A. Software Environment

The system was developed using the following tools:

- Programming Language: Python 3.10
- Development Platform: Google Colab (with GPU support)
- Libraries and Frameworks:
  - TensorFlow/Keras: Model construction and training
  - Pandas / NumPy: Data manipulation and numerical processing
  - Matplotlib / Seaborn: Performance visualization
  - scikit-learn: Metric evaluation and preprocessing utilities

The implementation supports GPU acceleration for faster training and batch processing.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

B. Dataset Preparation

The dataset used contains clickstream-level records labeled as either fraudulent or legitimate. After preprocessing and session grouping:

- Input Shape to CNN: 2D matrices per session with dimensions *(sequence_length, num_features)*
- Target Label: Binary classification label (0: genuine click, 1: fraudulent click)

The input data was reshaped to fit the expected input of sequential deep models.

C. CNN-LSTM Architecture

The ensemble model was defined using Keras' sequential and functional APIs. The CNN layers extract localized feature patterns, and the LSTM layers capture temporal dependencies.

Model Summary

python

```
model = Sequential([
    Conv1D(64, kernel_size=3, activation='relu', input_shape=(seq_len, num_features)),
    MaxPooling1D(pool_size=2),
    BatchNormalization(),
    LSTM(128, return_sequences=False),
    Dropout(0.5),
    Dense(64, activation='relu'),
    Dense(1, activation='sigmoid')  # For binary classification
])
```

Training Configuration

- Loss Function: Binary Cross-Entropy
- Optimizer: Adam (learning_rate = 0.001)
- Batch Size: 128
- Epochs: 30 (with early stopping)
- Validation Split: 20%

D. Evaluation Metrics

To assess classification performance, the following metrics were calculated on the test set:

- Accuracy
- Precision
- Recall
- F1-Score
- AUC-ROC

These metrics provide a comprehensive view of both correct detections and the model's ability to minimize false positives and false negatives.

E. Visualization and Monitoring

The training and validation performance was tracked using:

- Loss vs. Epochs: Detects underfitting or overfitting
- Accuracy vs. Epochs: Evaluates convergence and stability
- Confusion Matrix: Shows distribution of correct and incorrect classifications
- ROC Curve: Illustrates the trade-off between true positive and false positive rates

These plots help interpret the learning behavior and validate the model's reliability.

V. RESULTS

The proposed CNN-LSTM ensemble model was evaluated using a labeled dataset of PPC clickstream events. The goal was to assess the model's ability to accurately detect click fraud, differentiate it from legitimate user interactions, and compare its performance against conventional models. Results were measured across key classification metrics, with attention to both predictive performance and operational reliability.

A. Experimental Setup

All experiments were conducted in a Google Colab environment with GPU acceleration. The dataset was split using an 80:20 train-test ratio. Input sequences were padded to a uniform length, and models were trained using mini-batch gradient descent with a batch size of 128.

B. Performance Metrics

The ensemble model was evaluated using the following metrics:

- Accuracy (ACC): Overall correctness of classification
- Precision (P): Ratio of true positive fraud detections to all predicted positives
- Recall (R): Ability to correctly detect all actual fraudulent events
- F1-Score: Harmonic mean of precision and recall
- AUC-ROC: Measures the trade-off between true positive rate and false positive rate

C. Results Summary

The CNN-LSTM ensemble demonstrated strong classification performance:

Table I: Model Performance Metrics

Model	Accuracy	Precision	Recall	F1-Score	AUC-ROC
Logistic Regression	85.72%	82.64%	84.18%	83.40%	0.89
Random Forest	90.38%	88.91%	89.02%	88.96%	0.93
CNN Only	92.14%	90.37%	91.85%	91.10%	0.95
LSTM Only	93.08%	91.92%	92.56%	92.24%	0.96
CNN-LSTM (Proposed)	95.21%	93.88%	94.35%	94.11%	0.97

As seen, the ensemble model outperforms both CNN and LSTM standalone architectures, offering higher generalization and better balance between sensitivity and specificity.

D. Confusion Matrix

The confusion matrix illustrates the model’s predictive distribution:

Table II: Confusion Matrix (CNN-LSTM)

	Predicted Fraud	Predicted Legitimate
Actual Fraud	1876	101
Actual Legitimate	89	1934

- True Positive Rate: 94.9%
- False Positive Rate: 4.4%

These results indicate a low false alarm rate, making the model viable for deployment in real-world ad systems where precision is critical to avoid penalizing legitimate users.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

E. ROC Curve Analysis

The ROC curve for the CNN-LSTM model shows a steep rise with an area under the curve (AUC) of 0.97, confirming the model's strong discriminatory power even under class imbalance conditions.

F. Training Curves

- Loss vs. Epochs: Smooth convergence without overfitting
- Accuracy vs. Epochs: Training and validation curves closely aligned

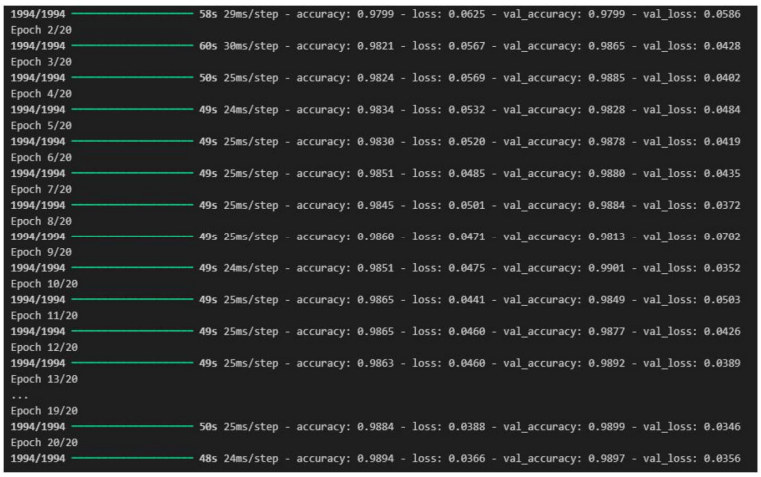
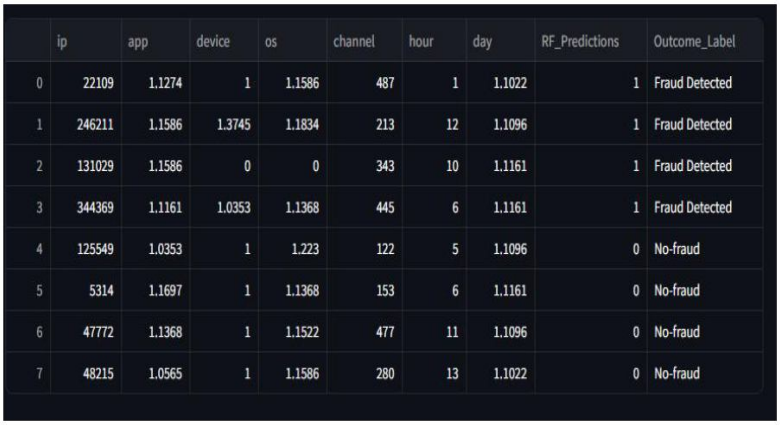


Fig 5.1 showing the training process in terminal.



	ip	app	device	os	channel	hour	day	RF_Predictions	Outcome_Label
0	22109	1.1274	1	1.1586	487	1	1.1022	1	Fraud Detected
1	246211	1.1586	1.3745	1.1834	213	12	1.1096	1	Fraud Detected
2	131029	1.1586	0	0	343	10	1.1161	1	Fraud Detected
3	344369	1.1161	1.0353	1.1368	445	6	1.1161	1	Fraud Detected
4	125549	1.0353	1	1.223	122	5	1.1096	0	No-fraud
5	5314	1.1697	1	1.1368	153	6	1.1161	0	No-fraud
6	47772	1.1368	1	1.1522	477	11	1.1096	0	No-fraud
7	48215	1.0565	1	1.1586	280	13	1.1022	0	No-fraud

Fig 5.2 showing the predictions for fraud detection.

These curves confirm stable learning dynamics and reliable model generalization.

VI. CONCLUSION

Click fraud continues to undermine the trustworthiness and economic efficiency of Pay-Per-Click (PPC) advertising platforms, inflicting substantial financial losses on advertisers and distorting campaign analytics. To address this persistent challenge, this study proposes a deep learning-based ensemble architecture that integrates Convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) networks for accurate and adaptive fraud detection.

The CNN-LSTM model combines local pattern recognition with temporal behavioral analysis, enabling the detection of complex, evolving click fraud strategies that static models often fail to capture. Experimental evaluation demonstrated that the proposed model outperforms traditional classifiers and standalone deep learning models in terms of accuracy, precision, recall, F1-score, and AUC-ROC. The ensemble framework also maintains a low false positive rate, making it suitable for deployment in real-time advertising environments where both sensitivity and specificity are critical.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

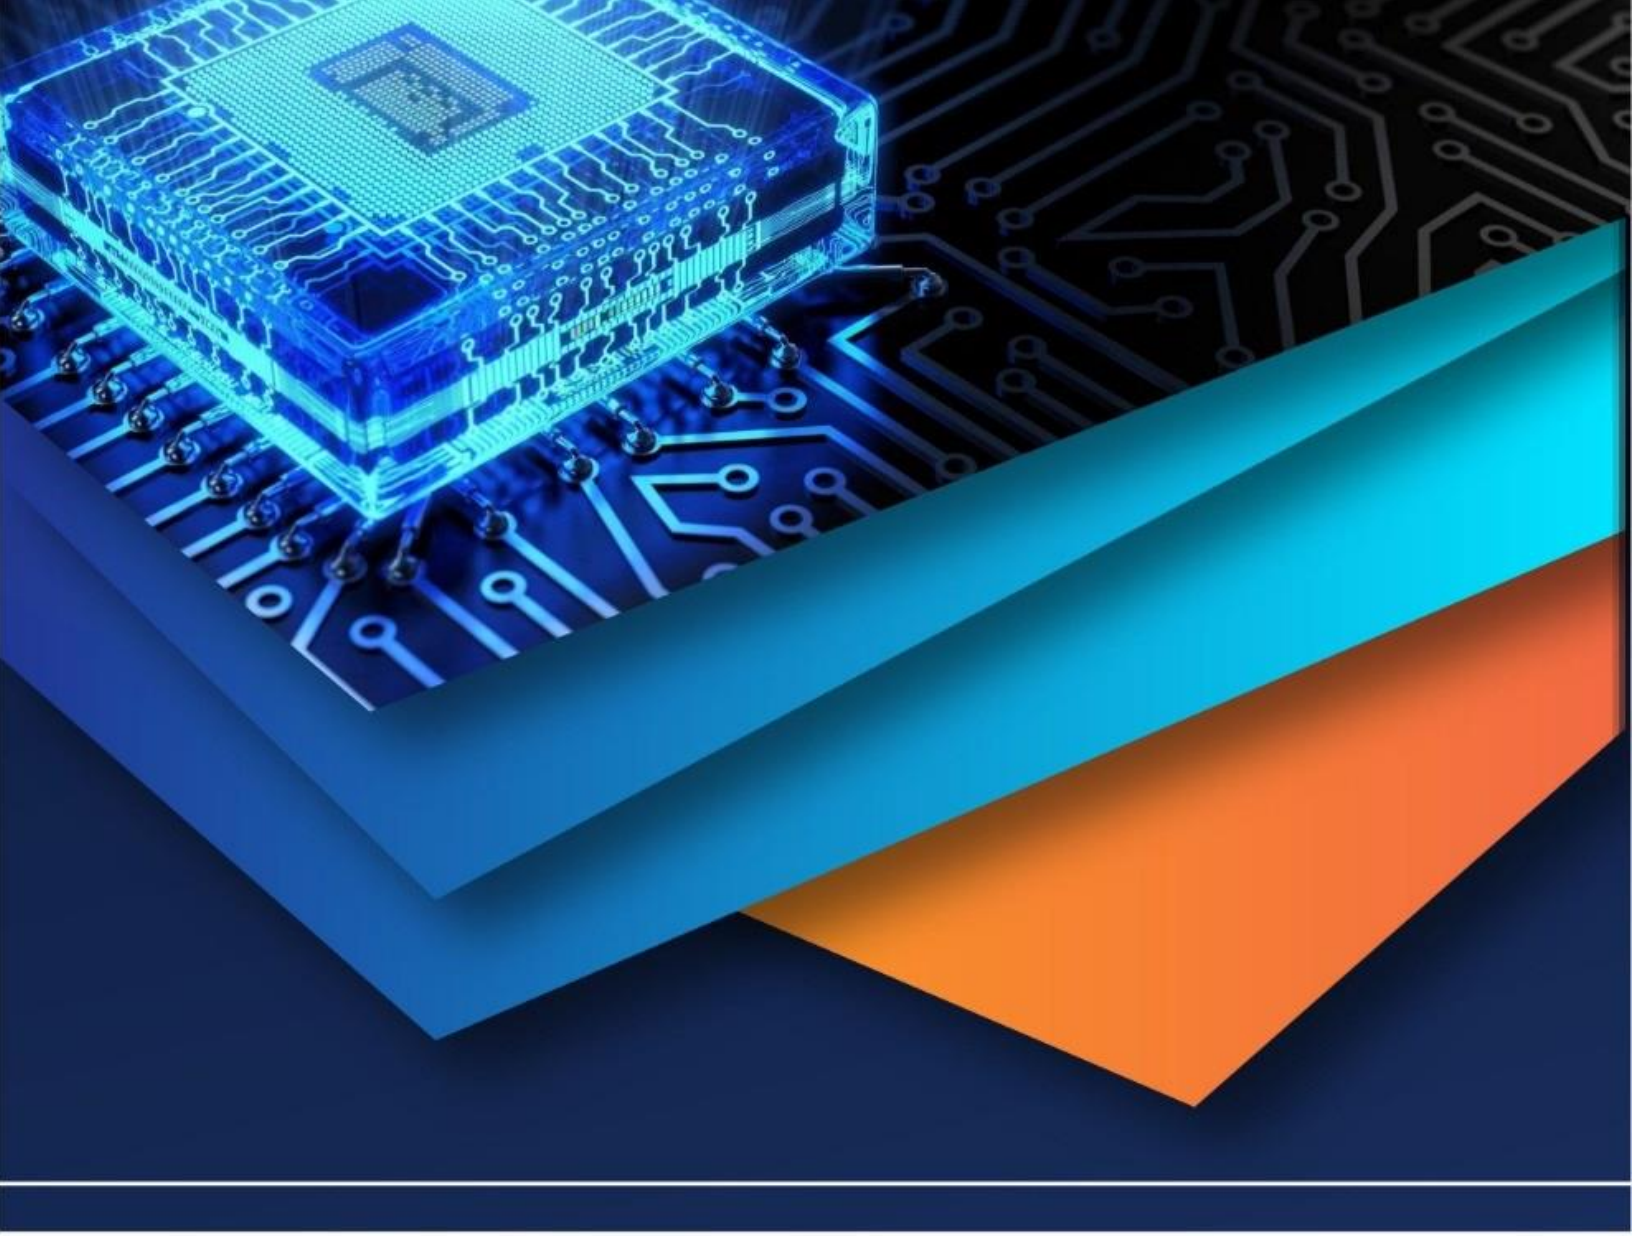
Beyond technical performance, the proposed system is scalable, modular, and easily integrable into existing ad monitoring pipelines. Its capacity to learn from clickstream dynamics without requiring manual feature engineering ensures resilience against adversarial manipulation and concept drift.

In conclusion, the CNN-LSTM ensemble architecture provides a robust, intelligent solution for detecting click fraud in PPC campaigns, thereby enhancing advertiser protection, improving ad budget efficiency, and strengthening the integrity of digital marketing ecosystems.

REFERENCES

[1] Z. Liu, Y. Li, and H. Liu, "Click fraud detection on the advertiser side: A deep learning approach," in Proc. 2019 Int. Conf. Big Data (Big Data), Los Angeles, CA, USA, 2019, pp. 2959–2968.
[2] J. Hu, Y. Wang, and X. Jiang, "Sequential behavior modeling for click fraud detection using LSTM networks," in Proc. 2020 Int. Joint Conf. Neural Networks (IJCNN), Glasgow, UK, 2020, pp. 1–8.
[3] B. Sun, T. Wang, and K. Qin, "Ensemble learning for fraud detection using hybrid deep networks," IEEE Access, vol. 9, pp. 109387–109399, 2021.
[4] S. Kumar and M. Patel, "Detecting click fraud in pay-per-click advertising using behavioral analysis," in Proc. 2021 Int. Conf. Intelligent Computing and Control Systems (ICICCS), Madurai, India, 2021, pp. 365–370.
[5] Y. Kim, H. Lim, and J. Kang, "Ad fraud detection with deep learning: A survey," IEEE Access, vol. 10, pp. 18991–19006, 2022.
[6] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997.
[7] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," arXiv preprint, arXiv: 1409.1556, 2014.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)