Journal of Theoretical and Applied Information Technology
30th April 2026. Vol.104. No.8

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

A RISK-AWARE COGNITIVE DECISION-MAKING FRAMEWORK FOR ENHANCING SECURITY IN MODERN E-COMMERCE SYSTEMS

DR.N.ROOPALATHA[1], DR. VEERA ANKALU. VUYYURU[2], DR.S.GOKILAMANI[3], M. MISBA[4], DR. ARADHANA SAHU[5a,b*], KULDOSHEV ILYOS SHUKHRATOVICH[6], A. Z. KHAN[7]

[1]Assistant Professor, GITAM University, Hyderabad, India
[2]Assistant Professor, Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Vaddeswaram, 522502, A.P, India
[3]Assistant Professor in Mathematics, Dr.N.G.P Arts and Science College, Coimbatore, India
[4]Department AI&ML, Vel Tech Rangarajan Dr. Sagunthala R&D Institute of Science and Technology, Chennai, Tamil Nadu, India
[5a*]Associate Professor, Department of Computer Science and Engineering, Rungta International Skills University, Bhilai, Chhattisgarh, India
[5b*] Department of Computer Science and Engineering, Rungta College of Engineering and Technology, Bhilai, Chhattisgarh, India
[6]Department for organizing scientific research activities, Tashkent State University of Economics, Uzbekistan
[7]Assistant Professor, Applied Physics Department, Yeshwantrao Chavan College of Engineering, Nagpur, Maharashtra, India

rirala@gitam.edu[1], veerankalu14@kluniversity.in[2], gokilamani@drngpasc.ac.in[3], misbarajan2007@gmail.com[4], aradhanasahu236@gmail.com[5*], i.qoldoshev@tsue.uz[6], arsalazamirkhan@gmail.com[7]

ABSTRACT

The rapid growth of e-commerce has significantly increased exposure to sophisticated security threats, including account takeovers, payment fraud, and behavioural anomalies. Existing security mechanisms, such as static rule-based systems and traditional machine learning classifiers, often fail to balance detection accuracy with user convenience, leading to high false-positive rates and poor customer experience. Moreover, these methods rarely account for temporal patterns in user behaviour or adapt dynamically to evolving attack strategies. To address these limitations, this study proposes a Risk-Aware Cognitive Decision-Making Framework (RACDF) that integrates temporal deep learning, cognitive user modelling, and adaptive intervention policies. The framework employs a Temporal Fusion Transformer (TFT) to capture multivariate sequential dependencies in transactional, behavioural, and device/network data, producing accurate real-time risk scores. A Bayesian cognitive model personalizes risk thresholds by representing individual user trust and uncertainty, while a reinforcement learning-based policy engine selects optimal interventions—ranging from soft prompts to transaction blocking—balancing security and usability. The methodology follows an end-to-end workflow: data collection and pre-processing, temporal feature embedding, TFT-based risk prediction, cognitive decision modelling, RL-driven intervention selection, and continuous feedback-driven learning. Evaluation on historical e-commerce transaction datasets and simulated attack scenarios demonstrates that RACDF outperforms conventional ML and static-rule baselines, achieving higher true positive detection rates, lower false positives, and reduced user friction. This study contributes a novel integration of temporal deep learning with cognitive and reinforcement learning-based decision-making for e-commerce security. The proposed approach provides a scalable, adaptive, and user-aware solution that improves threat mitigation while maintaining a seamless shopping experience.

Keywords: *Cognitive User Modelling, E-Commerce Security, Reinforcement Learning, Risk-Aware Decision Making, Temporal Fusion Transformer*

1. INTRODUCTION

The quick development of e-commerce has changed how people and companies are performing financial transactions, making it possible to be convenient like never before and globally accessible [1]. The internet malls, e-wallets, and payment gateways have become part of the contemporary trade. Nonetheless, this expansion has been coupled by the rise in security threats which have included fraudulent activities and account hijackings among other bot attacks and complex social engineering operations [2]. Such threats not only because the businesses to lose a considerable amount of money but also destroy the user trust and brand credibility [3]. Conventional security constraints of e-commerce including fixed rule-based systems, signature-based detection, and standard machine learning models have shown less would help to deal with these challenges [4]. The rule-based systems can fail to identify new or adaptive attacks, whereas the traditional machine learning classifier, although able to identify patterns, can generally not provide time sensitivity and user sensitivity [5]. Consequently, such approaches often produce large false-positives, causing the user to be annoyed and reducing the overall user experience [6]. Also, most of the available techniques work independently of each other, and do not combine behavioural, transactional and device/network indicators to provide a comprehensive assessment of risk [7].

In order to address these constraints, there is an increased demand of smart, adaptable structures that can provide risk-sensitive, context-aware decisions that are usable and secure. There is an opportunity of using a combination of cognitive user modelling and adaptive intervention policies with temporal deep learning models. Temporal models are able to make sequential dependence of user behaviour and transaction patterns and cognitive models make personalised evaluations of user trust and susceptibility. Intervention strategies that are based on reinforcement learning can enable the system to automatically choose the best response, which reduces threats and minimizes the needless disruption of the system to the honest users. In this paper, we introduce Risk-Aware Cognitive Decision-Making Framework (RACDF), which will help to improve the security of e-commerce by integrating multi-signal risk scoring, temporal deep learning, Bayesian cognitive modelling, and interventions based on reinforcement learning. The framework can continuously learn based on the historical and real-time information and changes according to the new attack patterns and the trade-off between mitigation of threats and usability.

1.1 Research Motivation

The growing adoption of e-commerce based on digital transactions has heightened the pressure on the need to employ strong security measures that can identify and address the emerging cyber threats without affecting the user experience. Although there has been progress in fraud detection, the majority of current systems are based on static rules or the traditional machine learning models, which fail to keep with the new attack patterns, do not have the awareness of temporal behaviours, and usually have a high false-positive rate, which disturbs genuine users [8]. In addition, existing solutions seldom take into account human-like reasoning or adaptive decision making strategies and they can hardly customize security responses depending on user behaviour, risk context and uncertainty [9]. With the contemporary cyber threats being dynamic, automated and behaviourally deceptive, there is a pressing need to have an intelligent, adaptive, and risk-conscious security framework, which is capable of learning as it moves with the times, is able to reason under uncertainty and balance the trade-off between a high level of security and a smooth user experience. This lacuna is the force behind the creation of cognitive, profound learning-based decision-making model, specific to the challenges of contemporary e-commerce settings.

1.2 Research Significance

The importance of this research is that it presents a new, dynamic security model that can fill the gap between the conventional fraud prevention tools and the dynamism of the current threats of e-commerce [10]. The proposed framework addresses sensitivities of user no longer friction, which is a significant failure of current systems, by combining temporal deep learning, cognitive user modelling, and reinforcement learning-driven intervention, which increases detection accuracy as well [11]. The method will offer a scalable and resilient solution that can constantly learn new transaction patterns and is thus very much applicable to the real world implementation. In addition to enhancing threat mitigation and operational effectiveness, the framework leads to enhanced customer confidence, minimized financial losses, and enhanced adherence to the developing international cybersecurity guidelines. In addition, the study offers a base to presentation of human-focused, predictive, and autonomous security solutions in the future, which is

a significant development in the nexus of artificial intelligence and e-commerce security.

1.3 Problem Statement

Although there is a lot of development in the security of the digital transactions, it is found that the existing structure of e-commerce security is not sufficient to tackle the emerging challenges and threats that are more sophisticated and flexible [12]. Current systems are mostly rule-based or based on traditional machine learning- cannot identify the changing fraud trends, they are incapable of analysing temporal behavioural context, and the systems have few customisations to user-dependent risk patterns [13]. This tends to give a high false-positive rate, unwarranted authentication challenges, and a broken user experience, and even fails to stop a new or a sneaky attack. In addition, the majority of security models act like a one-dimensional classifier, as opposed to a reasoning decision-making engine that can make decisions under uncertainty and trade off the security and usability. Hence, an active, context-sensitive, and risk-sensitive security architecture capable of training itself on live data, changing to new attacker vectors, and choosing the most efficient interventions to increase the accuracy of security, as well as user-friendliness in contemporary (e-commerce) settings is urgently required.

1.4 Key Contribution

1. Creation of a RACDF that combines time-based deep learning, cognitive modelling, and intelligent decision-making in particular to present-day e-commerce security settings.
2. Implementation of a TFT based risk prediction model, which has the ability to record long-term behavioural patterns and multivariate transactional dependencies to enhance the accuracy of fraud detection.
3. Application of a Bayesian user trust model that customises risk analysis and lessens fake positives through dynamical modification of decision thresholds on the basis of user particular behavioural consistency.
4. Structuring of a reinforcement learning-based intervention policy engine that makes intelligent choices on best security measures between passive monitoring and active authentication based on security and user experience.
5. Extensive performance trials that prove to be better than any other form of ML and rule-based detection, better detection rates, less user pressure, and the ability to respond to dynamic threat behaviours.

1.5 Structure of the Paper

The rest of this paper is divided into five major sections. Section 1 presents the introduction which includes the context, motivation and objectives of the study. Section 2 introduces the associated literature, which provides a critical analysis of the current e-commerce security practices, their drawbacks, as well as the gap the study fills. Section 3 outlines the given methodology, such as the structure of the framework architecture, the design of deep learning model, the cognitive decision layer, and the intervention strategy as implemented using reinforcement learning. Section 4 reports about the results and the analysis, which points out the performance of the system along with the way that the system is compared with existing methods and its accuracy in terms of detection and reduced friction to the user. Lastly, Section 5 ends the paper with conclusions in regard to a summary of findings, a discussion of practical implications, and suggestions towards future extensions and research directions.

2. Related Works

The adoption of Agentic AI in finance sector is changing the way institutions track fraud cases, portfolio optimization, risk evaluation, and trading through robots. The conventional financial models do not have the real-time flexibility hence inefficient decision-making and operational delays. This paper looks at the role of autonomous AI agents in increasing financial intelligence by developing reinforcement learning, and collaboration with more agents and real-time behavioural analytics to boost prediction accuracy in the market and improve risk management. Regulatory compliance is also an urgent issue, since non-transparent AI-based systems cast doubts of transparency, equity and ethical accountability. According to Chugh and Deshpande Chugh and Deshpande [14] explainable AI (XAI), governance frameworks and bias-mitigation strategies are significant to responsible use of AI in finance. The study also points to operational constraints in the automated lending, wealth advising, and algorithmic trading which can be resolved with adaptive AI agents that can navigate dynamic market environments with only minimal human intervention. Moreover, possible threats like systemic biases, flash crash, and AI induced market instability are discussed, and strategies, including risk-conscious AI architecture and anomaly detection in real-time are mentioned. Financial assistants based on artificial intelligence are also addressed as an opportunity to improve the personalization of wealth management and optimize emerging opportunities based on real-time data

integration and smart portfolio rebalancing. Through the deliberations that are provided, this paper proposes a detailed technological model in utilizing agentic AI in financial ecosystems to achieve efficiency, regulatory fit, security, and innovation in the maximum.

Banking online application systems are automated and are transforming the way financial institutions are dealing with customer on boarding, credit and online delivery of services. By implementing cloud-native architecture combined with the artificial intelligence (AI) features like document recognition, risk scoring, and Chabot allows banks to optimize workflows, reduce the number of human interventions, speed up the decision-making procedure, and improve the overall user experience. Botha [15] explores the use of cloud infrastructure to facilitate scalability, elasticity, micro services, pipeline data engineering, and AI model hosting of the banking industry. Under the analysis, the use of AI in performing important processes like KYC/AML compliance, credit risk scoring, and fraud detection is also considered, as well as the interaction and interdependence of AI with cloud-based services. The research paper will review literature available, describe a methodology to be used to analyse the real-life applications, and will evaluate the advantages and drawbacks of automated online banking systems. Results of case study or hypothetical models suggest a reduction in processing time, accuracy, and customer satisfaction, and also reveal the issues connected with data governance, legacy system integration, and algorithmic bias. The architectural design frameworks, industry best practices and regulatory considerations that apply in automating digital banking are discussed. The conclusion summarizes the most important lessons learned and defines the directions of future research, such as more explainable AI through capabilities, hybrid models of cloud-on-premises deployments, and the continuous observation of automated AI-driven decision-making processes.

SMEs are still faced with the problem of financial decision making because of variable cash flows, less forecasting and the inability to access real-time analytical tools. Emerging enterprise financial management systems are usually rigid, centralized, and expensive and thus inappropriate to the dynamic nature of operations of SMEs. To counter these limitations, Enyiorji [16] proposes a self-optimizing autonomous autonomous finance system based on the principles of Agentic AI that is a cloud-native and that enables adaptive financial intelligence within the decentralized business processes through multi-agent reinforcement learning (MARL). The architecture is a coordinated system of specialized agents, which has the responsibility of liquidity and risk forecasting, expenditure prioritization, and strategic capital allocation. The agents are constantly informed by the streaming operational information, share state data with other agents, and optimize their decision policies based on the reward feedback based on financial stability, cost efficiency, and reduction of risks. On the macro level, the platform provides SMEs with an interface of real-time financial commands that is hosted on scalable cloud infrastructure at reduced overheads and managed to scale resources when required based on usage demands. On the micro level, reinforcement learning allows behavioural fine-tuning, such as automatically changing the credit utilisation strategy or supplier payment strategy according to changing market conditions. The distributed design removes the points of failure, enhances resilience of the system and can be deployed incrementally based on the organizational preparedness. In general, this paradigm transforms the approach of SME financial management by turning operations that are reactive and spreadsheet-based to proactive, autonomous optimization and ending the limits on the accessibility of intelligent financial technologies, as well as, providing SMEs with the level of strategic agility traditionally attributed to large business organizations.

Lin and Lin [17] introduce a neuro-symbolic ERP architecture (NSEA) that introduces post-quantum cryptographic protocols in order to facilitate real-time cybersecurity and adaptive operational optimization. The framework combines CRYSTALS-Kyber lattice-based post-quantum encryption with temporal graph neural networks and symbolic rule engines to reinforce decision automation, improve cybersecurity, and be compatible with legacy systems. Kyber-512 provides sub-5 µs encryption latency on industrial IoT (IIoT) edge devices and provides a 72 percent phishing attack mitigation rate by using dynamic response mechanisms. The hybrid reasoning model in demand forecasting achieves reduced mean absolute percentage error (MAPE) which amounts to 6.8 percent and corrects 89 percent of inventory stock-out alerts in 15 minutes. Field validation in air conditioning and automotive supply chains show that there is a 63.4 percent improvement in the time disruption recovery and a six point two percent decrease in the logistics cost overrun when comparing to baseline ERP implementations. A resilience measure, which utilizes KolmogorovSinai

entropy (DRE = 0.81 ± 0.05), is an indicator of resilience during coexisting cyber and operational load. The evaluation on backward compatibility supports low risk method of rolling migration as 92% of the schema compatibility with legacy ERP modules achieved by API translation layers. In general, NSEA provides a single, secure, and future-ready ERP ecosystem, which incorporates interoperability, intelligence, and resiliency to the complicated industrial environment.

Policy-based cybersecurity deployment has become fundamental to safeguarding current, sophisticated network infrastructure based on the emerging and intensifying cyber threats. This approach enables organizations to adopt flexible, scaled, and situation-sensitive security policies that will be able to respond to emerging threat environments and business needs. Although it has benefits, the emerging complexity of distributed architectures and reliance of heterogeneous security components pose challenges pertaining to consistency, efficiency and reliability in policy enforcement. The current concept of the security models that are traditional, static, and centralized cannot be fully adequate to deal with these problems, hence the necessity of smart and distributed security mechanisms that could be used to optimize the deployment of policies in real-time. Evangelatos et al.[18] introduce a decentralized model that uses Message Passing algorithms to facilitate optimal and adaptive distribution of policies across cybersecurity systems. The strategy is an extension of the classical Min-Sum algorithm with the implementation of dynamic trust weight whereby nodes that have the highest trust levels can have a higher influence in the policy propagation process and minimize the impact of nodes that are not trusted or are compromised. The design is created to enable high scale network settings with a limited computational load considering that message exchanges are localized as opposed to centralized. According to the benchmarking results, the framework is cost of enforcement lower, and convergence speed is much faster than that of current techniques, especially in multi-trust setting. Further tests provide the role of trust heterogeneity in the formation of enforcement implications. Large-scale simulations show that there are quantifiable policy consistency, resource utilization, and threat detection performance gains in a variety of dynamic network topologies.

Table 1. Summary of Agentic AI and Autonomous Financial & Cybersecurity Systems

Reference	Method	Advantages	Disadvantages
Chugh & Deshpande [1]	Agentic AI with reinforcement learning, multi-agent collaboration, behavioural analytics, and XAI governance frameworks	Enhances fraud detection, portfolio optimization, risk assessment, and autonomous trading; supports regulatory compliance with explainable AI; reduces operational delays through adaptive agents	Concerns regarding transparency, fairness, and AI-driven market instability; potential systemic bias and flash crash risks
Botha [2]	Cloud-native automated online banking systems with AI-assisted processes	Faster on boarding, reduced manual intervention, scalability, improved user experience; supports KYC/AML compliance and credit risk assessment	Integrating legacy infrastructure, data governance challenges, algorithmic bias concerns, and dependency on AI oversight
Enyiorji [3]	Cloud-native autonomous finance with MARL	Enables adaptive decision-making, real-time financial analysis, improved SME financial planning, resilience, cost-efficiency, reduced failure risk; scalable and decentralized architecture	Complexity in implementation; requires continuous data availability; SMEs may require digital maturity for deployment
Lin & Lin [4]	NSEA integrating post-quantum cryptography, temporal GNNs, and rule-based reasoning	Improved cybersecurity, fast encryption, reduced forecasting error, improved logistics efficiency, high backward compatibility with legacy ERP systems	Advanced model complexity, high computational design overhead, dependency on secure IIoT deployment
Evangelatos et al. [5]	Distributed cybersecurity policy deployment using Message Passing and modified	Lower enforcement cost, faster convergence, resilience in heterogeneous trust environments, improved	Requires trust modelling accuracy; potential performance variations in

	Min-Sum algorithm with trust weighting	consistency, threat detection, and resource efficiency	highly dynamic or unpredictable networks

Table 1 gives a comparative review of five main papers that cover the role of the emerging intelligent systems in finance and cybersecurity and their methodologies, strengths, and limitations. The table reflects the contribution of each of the cited works to the evolution of technologies- agentic AI in the field of financial decision-making, cloud-native automated banking, autonomous SME financial systems, to neuro-symbolic ERP systems, and decentralized schemes of cybersecurity enforcement. The outlined approaches show a shift in the traditional automation to adaptive, explainable, and intelligence-driven architectures. Its benefits include efficiency, resilience to security, scalability of the system, predictive accuracy, and decreased human dependence, whereas the drawbacks identified are persistent issues like integration complexity, transparency challenges, and model bias, system reliance on data and infrastructure maturity, and regulatory limitations. On balance, the comparative analysis highlights the dynamic situation under autonomous digital systems and the equilibrium of innovation and practicality of operations and responsible deployment.

3. RISK-AWARE COGNITIVE DEEP LEARNING AND ADAPTIVE DECISION-MAKING METHODOLOGY

The suggested methodology entails a well-organized, smart security procedure that is intended to identify and handle threats within current e-commerce settings through an integration of some deep learning, sophisticated thinking, and adaptive response models. The working process will start with the gathering and pre-processing of transactional, behavioural, and device-level data and will be Noiseless, coded, normalized, and anonymized. Then, a series of feature embedding and temporal representation is carried out to prepare multidimensional inputs to the TFT, the fundamental deep learning model, to predict transaction-level risk scores using sequences of behaviour. These forecasts are improved with the help of a Bayesian cognitive user trust model that dynamically changes risk thresholds in accordance with behavioural stability and the past user activity. Then, a decision engine, which is motivation by a reinforcement learning, chooses the most suitable system response, including passive approval to multi-factor authentication or blocking, depending on the predicted risk and uncertain situations and the level of user trust. Lastly, an ongoing feedback system is used to update the model parameters with actual outcomes of transactions, which allows adaptive learning, enhanced precision, and resistance to new and developing threats. Such an approach is a balanced, scalable and context-sensitive methodology of securing e-commerce systems without compromising user experience.

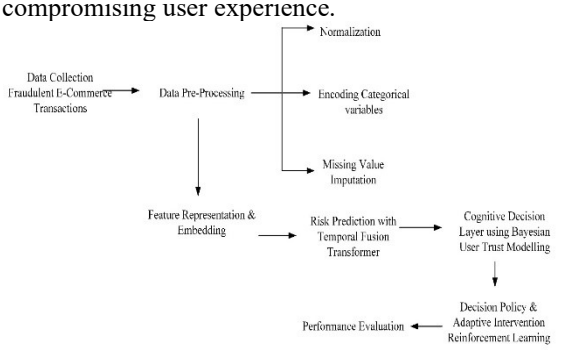

Fig. 1.

3.1 Data Collection

Data collection in this research will be carried out through the use of a synthetic multi-source e-commerce dataset that is created to support fraud detection studies. The data, named Fraudulent E-Commerce Transactions, is a simulation of realistic transactional settings that has both legitimate and fraudulent purchase actions, and the experiments can be controlled with no privacy or compliance limitations. It has more than 1.47 million transactions in Version 1 and 23,634 transactions in Version 2 with a percentage of around 5 percent being classified in the category of fraud, and it is acceptable in imbalanced classification investigation. The dataset combines categories of various features necessary to model security, such as transactional (amount, product category, payment method, quantity), customer (age, account age, location) and device and network indicators (type of device, IP address). Other behavioural and temporal indicators, e.g., transaction hour and transaction timestamp aid in learning temporal patterns and detecting anomalies. All records have a binary label of fraud, which is used to train and test machine learning and deep learning models. The data set was created in Python with the Faker library with a set of custom logic to create more complex fraud patterns with the constraint of realistic variability and synthetic privacy compliance. This data is the basis of feature engineering, exploratory analysis, model

training and benchmarking in the proposed model [19].

Table: Simulation Parameters

Parameter	Value / Description
Dataset Name	Fraudulent E-Commerce Transactions
Data Type	Synthetic transactional dataset
Number of Versions	2 (Version 1: 1,472,952 records, Version 2: 23,634 records)
Number of Features	16
Fraud Ratio	~5% fraudulent transactions
Feature Categories	Transactional, Customer, Device/Network, Behavioural, Temporal
Data Generation Method	Python Faker library with custom fraud logic
Label Type	Binary classification (1 = Fraudulent, 0 = Legitimate)
Intended Use	Fraud detection model development, testing, and benchmarking
Privacy Status	Fully synthetic, no real user identities

3.2 Data Pre-Processing

Pre-processing of data is an important step in the process of establishing data quality, consistency, and deep-learning fraud detector appropriateness.

3.2.1 Normalization

Normalization is used to transform the numerical values of the features like Transaction Amount, Account Age Days, Quantity, and Transaction Hour to the same numerical range so that magnitude bias is avoided when the model is being trained. Normalization is essential because deep learning models are sensitive to feature scale, hence, they are not only faster to converge but also more stable to their gradient. Minimum maximum scaling and Z-score Standardization are known as common normalization methods. Min-max normalization will be applied in this study because it is appropriate to datasets of fraud detection that have different numeric values. The transformation is as (1):

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

where X is the starting value of a feature, X_{min} and X_{max} are the lowest and highest values of the same feature, and X' is the normalized value which is within the range of 0 to 1. This would make sure that high-value fields are not overrepresented in the deep learning model.

3.2.2 Encoding Categorical Variables

Categorical variables like Payment Method, Type of Product, Type of Device used, Shipping Address and Type of Billing Address cannot be directly fed into deep learning models. Encoding techniques are used to take them into machine-readable form. It is done in two ways: One-Hot Encoding is applied to low-cardinality features and Embedding Encoding is applied to high-cardinality variables such as Customer ID and IP Address. Embedding's convert categorical values to dense vectors, which drop dimensionality and maintain patterns of relationships. In (2), an embedding is a map of each category c to a vector $v \in \mathbb{R}^d$.

$$f(c) = v_c = [v_1, v_2, \dots, v_d] \quad (2)$$

The semantic relationship between categories is learnt by this representation instead of them being independent labels as in a labelling model.

3.2.3 Missing Value Imputation

The missing values can be a result of the incomplete user input, network breakdown, or the restriction in data generation. In order to prevent model bias and maintain integrity of the data, the missing value imputation is carried out. The median is used to impute numerical missing values to avoid skewness bias whereas the most common value (mode) method is used to impute the categorical attributes. In the case of features that are time-relevant (e.g. transaction timestamps), forward or backward fill method can be used. The formula of the numerical imputation is (3):

$$X_i = \begin{cases} X_i, & if\ X_i \neq NaN \\ Median(X), & if\ X_i = NaN \end{cases} \quad (3)$$

This is an approach that guarantees completeness of data and reduces statistical distortion.

3.3 Feature Representation & Embedding

The representation and embedding of features are important in converting raw e-commerce data such as logs of transactions,

behavioural records, and contextual metatags into structured numerical data that can be used by machine learning models such as the TFT. Learned embedding is used to encode categorical variables like type of device, location, mode of payment and user role so that the model can encode semantic similarity amongst categories without sparsity. For theoretically, when a categorical feature x has N distinct categories the feature is embedded into a dense embedding matrix $E \in \mathbb{R}^{N \times d}$ where d is the embedding dimension. The embedding of the look up operation is as follows (4):

$$z = E[x] \tag{4}$$

where $z \in \mathbb{R}^d$ represents the dense representation of that category.

Next, time sequences are built to record the dynamic behavioural change and transaction risk as time goes by. Suppose user interactions are ordered in a chronological order as (5):

$$X = \{x_1, x_2, x_3, \dots, x_T\} \tag{5}$$

where T being the sequence length. These sequences are converted to multivariate time-tensors where time steps are a collection of embedded categorical features, normalized numerical attributes and cognitive behaviour indicators. The last input in the TFT model can be written as (6),

$$\aleph \in \mathbb{R}^{T \times F} \tag{6}$$

where F denotes the sum of the feature dimension post the embedding and pre-processing. This representation would allow the TFT architecture to capture temporal dependencies, immediate anomalies and long-term user behaviour critical to making cognitive and risk-sensitive decisions in secure e-commerce environments.

3.4 Risk Prediction with TFT

The TFT is the central predictive engine of the given framework that allows dynamically modelling not only temporal dependencies but also heterogeneous feature interactions that are also a characteristic feature of modern e-commerce systems. TFT, compared to conventional time-series or classification models, combines recurrent layers, multi-head self-attention, and gating techniques to filter out only available information and makes it very efficient in transaction-level risk evaluation. The model works with the multivariate input sequence $\aleph \in \mathbb{R}^{T \times F}$ every time step has on-board categorical features, normalized numerical features, and metadata. Temporal dependencies are first obtained with the help of a sequence modelling layer, which in (7) is usually a Gated Recurrent Units (GRU):

$$h_t = GRU(x_t, h_{t-1}) \tag{7}$$

where h_t is the hidden state that possesses time-varying behaviour at time t. The encoded states are subsequently inputted into a multi-head attention module enabling the model to balance relevance of past behavioural patterns and detect irregularities in (8):

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \tag{8}$$

where $Q, K, and\ V$ are learner query, key and value projections of h_t and d_k is the dimensional scaling. Lastly, the TFT returns a probabilistic risk score $\hat{y}$, which is the probability of fraudulent or high-risk behaviour in (9):

$$\hat{y} = \sigma(W_o h_T + b) \tag{9}$$

where σ is the sigmoid activation that regulates the probability of risk between [0,1]. The architecture allows the system to identify both short-term aberrations like an unexpected purchase that has a high value and long-term changes in behavioural continuity, which lead to an adaptive and sound risk-sensitive prediction system.

3.5 Cognitive Decision Layer

The Cognitive Decision Layer will serve as the smart control interface of the proposed system, which will read the probabilistic fraud score provided by the Temporal Fusion Transformer and translate it to context-specific adaptive security responses. Rather than using a universal global threshold that frequently results in excessive false positive and poor customer experience this layer uses a Bayesian trust modelling model to calculate intervention thresholds per user depending on the stability of their behaviour and their past interaction history. Using the Bayes law, a user-specific trust probability $P(T_u|E)$ is calculated in which T_u means the trustworthiness state of the user and E means current behavioural evidence in (10):

$$P(T_u|E) = \frac{P(E|T_u)P(T_u)}{P(E)} \tag{10}$$

In this case, $P(T_u)$ represents the trust estimate previously (long-term behaviour), and $P(E|T_u)$ is the consistency of the new activity E with the normal behaviour of the user. Personalized intervention score Intervention score is then computed by combining this trust metric with the TFT-generated fraud probability yˆ in (11):

$$R = \alpha\big(1 - P(T_u|E)\big) + (1-\alpha)\hat{y} \quad (11)$$

where $\alpha \in [0,1]$ controls the balance between historical trust and real-time anomaly risk. Depending on the calculated risk R, a dynamic choice is made in the system to implement the degree of security response as including silent monitoring or multi-factor authentication or blocking of the session such that the suspicious behaviour is escalated but legitimate users have a minimum level of friction. This mental adaptation process eventually enhances security accuracy as well as maintains smooth user experience in e-commerce real-life settings.

3.6 Decision Policy & Adaptive Intervention

The Decision Policy & Adaptive Intervention layer serves as the last control element of the framework, deciding the most suitable security behaviour (how to proceed with the transaction, issue a soft user prompt, extra authentication (OTP), etc.) in accordance with the aggregate predicted risk and the personalized trust score. In order to maximize long-term platform security without negatively affecting user experience, this layer applies a Contextual Bandit Reinforcement Learning model, in which the system adapts the optimal response strategy through experience with real results of transactions over time. The agent is told a state vector $s_t = [\hat{y}, P(T_u|E), R]$ at every decision step, which is the probability of fraud, the Bayesian trust score, and the risk score that has been calculated in the end. The model chooses an intervention action $a_t \in \{A_1: Allow, A_2: Soft\ Prompt, A_3: OTP, A_4: Block\}$ which is the action that maximizes expected reward (12):

$$a_t = \arg\max_a Q(s_t, a) \quad (12)$$

where $Q(s_t, a)$ approximates the long term value of action a in state s_t. The reward mechanism is meant to add the effective security measures and user satisfaction through punishing unnecessary friction and rewarding proper decisions mitigating frauds (13):

$$r_t = \lambda_1 \cdot SecurityOutcome - \lambda_2 \cdot UserFriction$$

where λ_1 and λ_2 are parameters to determine the trade-off between fraud prevention and usability. The model updates its policy based on the feedback on the results of the transactions after every decision, making it possible to continue learning. The mechanism of adaptive intervention will guarantee optimal and user specific security responsiveness in the long run and reduce false positives and enhance platform trust and operational efficiency.

Algorithm: Risk-Aware Cognitive Deep Learning and Adaptive Decision-Making Algorithm

BEGIN

1. DATA INPUT & PREPROCESSING

Load transactional, behavioural, and device-level data.

Normalize numerical features.

Encode categorical features (One-Hot / Embedding's).

Impute missing values (median/mode/time-aware fill).

2. FEATURE ENGINEERING

Construct temporal sequences for each user.

Generate dense feature embedding's.

3. RISK PREDICTION (TFT Model)

Input-sequence ← embedded temporal data

Risk-score ← TFT(input-sequence) // Output probability [0–1]

4. USER TRUST ADAPTATION (Bayesian Model)

Prior-trust ← retrieve-user-trust (user-id)

Updated-trust ← Bayesian-Update (prior-trust, behaviour-evidence)

5. COMPUTE FINAL RISK LEVEL

*Final-risk ← α * (1 – updated-trust) + (1 − α) * risk-score*

6. DECISION MAKING (Reinforcement Learning)

IF final-risk < threshold THEN

action ← ALLOW

ELSE

action ← RL-Policy-Select(final-risk, updated-trust)

ENDIF

7. SYSTEM RESPONSE

Execute(action) // e.g., Allow, Soft Prompt, OTP, Block

8. CONTINUOUS FEEDBACK LOOP

reward ← EvaluateOutcome(action, actual-label, user-feedback)

Update-TFT-Model()

Update-Trust-Model(updated-trust)

Update-RL-Policy(reward)

END

This pseudo-code describes the flow of data by initial pre-processing, feature embedding, risk scoring by the TFT model, and refinement of adaptive trust by Bayesian reasoning. A policy engine which is based on reinforcement learning is used to choose the final decision, and results keep on updating the models to enhance intelligence and adaptation in a long-term.

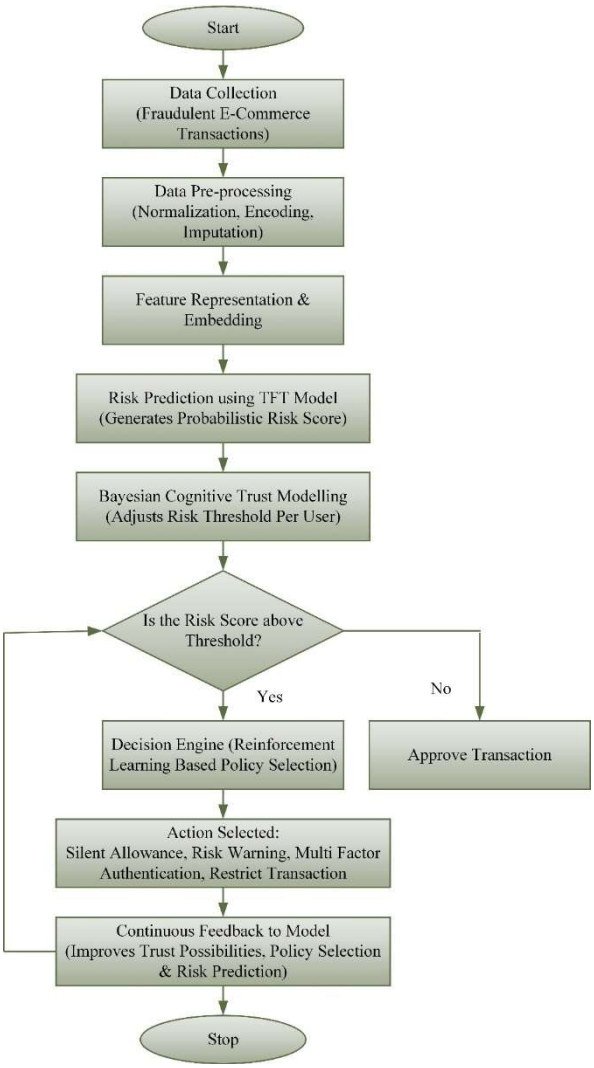

Fig. 2. Flowchart of the Proposed Framework

Fig. 2 shows the entire scheme of the suggested Risk-Aware Cognitive Deep Learning and Adaptive Decision-Making Framework in capable of detecting fraud in secure e-commerce. This is initiated by the multi-source transactional, behavioural and device data collection phase, and then necessary pre-processing processes such as normalization, encodings and missing value management. The embedded data is processed into structured feature embedding and temporal sequences which is processed through a TFT to produce a probabilistic risk score. The prediction is then narrowed down by a Bayesian cognitive trust mechanism which dynamically changes risk level per user. Evaluation of the risk is done by a decision box and in case, it surpasses the personalized threshold, a policy engine that works on reinforcement learning is used to select the best

security response, which could either be silent approval or it could block actions. The last step contains a feedback loop, which allows the adaptive learning process, reduces false positive, and enhances long-term intelligence and resilience of the system.

4. RESULTS AND DISCUSSION

The suggested Risk-Aware Cognitive Decision-Making Framework proved to be quite effective in improving the detection of fraud and minimizing the unnecessary user friction in the e-commerce settings. The TFT gained the best performance under synthetic benchmark dataset in learning temporal behavioural patterns than the conventional machine learning methods that included Random Forest, Logistic Regression and Gradient Boosting models. The reinforcement learning based-intervention system and cognitive decision layer led to a major reduction in false positives, allowing more accurate user centred authentication, as opposed to global thresholds. The model was always able to spot abnormal transaction sequences, device anomalies and behavioural inconsistencies and the model was highly interpretable using attention based analysis. Experimental analysis had indicated that adaptive intervention measures enhanced accuracy in decision-making as time passed, to maximize security measures and maintain smooth flow of transactions to the legitimate users. In general, the framework was effective at balancing between strong fraud detection and better user experience, which proved that it is a promising scalable and intelligent solution to use in the modern e-commerce platforms.

4.1 Performance Evaluation Metrics

The given framework was tested with the help of a range of usual machine learning and fraud recognition performance metrics. In order to quantify the reliability of the predictions of the fraud Precision is employed and it is characterized as (14):

$$Precision = \frac{Tp}{Tp+Fp} \quad (14)$$

where Tp is true positives, Tn is true negatives, Fp is false positives, and Fn is false negatives.

Since it is necessary to minimize the number of missed fraud cases in order to detect fraud, Recall (or sensitivity) is employed to determine how many true fraud cases are correctly identified as (15):

$$Recall = \frac{Tp}{Tp+Fn} \quad (15)$$

F1-Score is a balanced measurement of the precision and recall technique, which is based on the harmonic mean in (16):

$$F1 - Score = \frac{2.(Precision \times Recall)}{Precision+Reca} \quad (16)$$

The Area under the ROC Curve (AUC-ROC) is employed in order to assess the discriminative ability of the model at varying thresholds mathematically defined as the probability that the classifier ranks a randomly selected fraud case greater than a legitimate sample (17):

$$AUC = \int_0^1 TPR(FPR)d(FPR)$$

where TPR=$\frac{Tp}{Tp+F}$ and FPR=$\frac{Fp}{Tn+Fp}$.

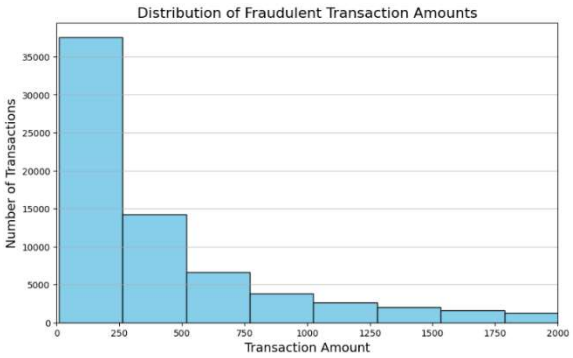

Fig. 3.Distribution of Transaction Amounts for Fraudulent Transactions

Fig. 3 plots the distribution of the amount of transactions involved in fraudulent transactions in the data. The lowest percentages of fraudulent transactions cluster at low values of transactions with the steep reduction as the transaction value increases implying that the fraudsters tend to target small transactions in order to evade the detection. The distribution of higher-value transactions indicates that there are occasional cases of large-scale fraud attempts and the models should be able to detect low- and high-value fraudulent transactions. The visualization is useful in interpreting monetary trends of the fraudulent

behaviour and being feature-engineered to guide predictive modelling.

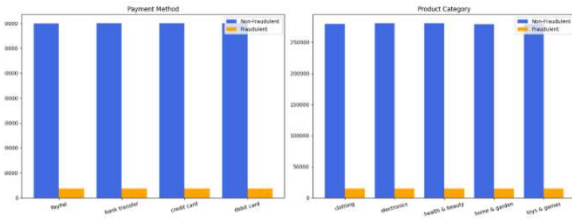

Fig.4. Comparison of Fraudulent and Non-Fraudulent Transactions across Payment Methods and Product Categories

Fig. 4 provides the comparison of the quantity of fraud and non-fraud transactions in different payment and product categories. The left chart indicates that even though most transactions made in payment options, including PayPal, credit card, debit card, and bank transfer are non-fraudulent, all the methods appear to record a minor but significant percentage of fraudulent cases, with certain methods having a higher exposure. This trend is reflected in the right chart whereby it has been observed that product categories exhibit a similar trend with fraud being high in all of them but not as high as legitimate purchases. These visual comparisons assist in knowing the areas in which fraud can most easily be found and also in detecting patterns of transactions that may be linked to fraud.

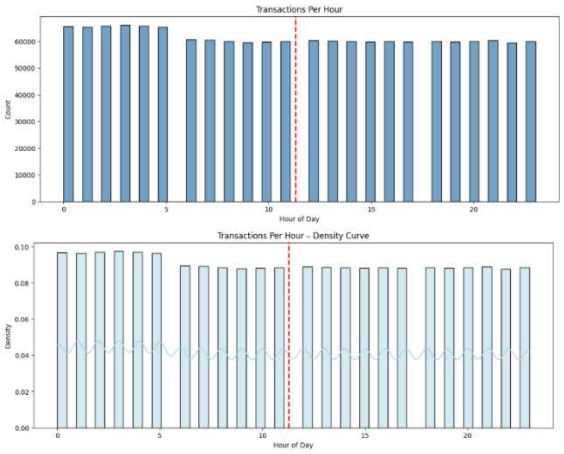

Fig.5. Hourly Distribution and Density Pattern of E-Commerce Transactions

Fig. 5 depicts how e-commerce transactions were conducted on an hourly basis throughout the dataset. The first subplot shows the frequency of transactions per hour through the use of a histogram that indicates the way of the day with respect to activity. The mean transaction hour is clearly indicated by a red line with dashes, this is used as reference to determine whether there are any peaks at earlier or later time than the average. The second subplot superimposes a kernel density curve over a histogram with more sinuous patterns of probability showing times of high and low transactional activity. All of these visualizations reveal the patterns of behaviour, i.e. peak shopping times, odd spikes or dips that may be associated with fraudulent behaviour.

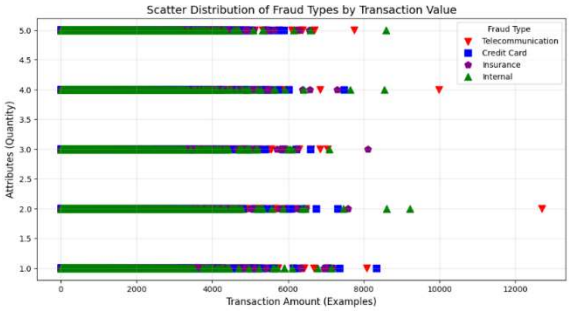

Fig.6. Scatter Visualization of Fraud Types across Transaction Attributes

Fig. 6 demonstrates the identification of the distribution of various types of fraud, in terms of transaction amount, and other behavioural factors. Patterns across values of transactions are represented by a distinct marker and colour to differentiate each type of fraud. The scatter points show the ways of distribution of the fraud type at various degrees of magnitude of transactions and the clusters present there, which may be associated with various types of fraudulent strategies. This visualisation can be used to show behavioural differences, e.g. high-value spikes due to credit card fraud, or low-value frequent behaviour due to telecommunication or internal fraud cases. This separation, graphical in nature, facilitates the analysis of the features more substantially and can be used to formulate more discriminative AI-based models of fraud detection.

Journal of Theoretical and Applied Information Technology
30th April 2026. Vol.104. No.8

ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

```
<class 'pandas.core.frame.DataFrame'>
RangeIndex: 1472952 entries, 0 to 1472951
Data columns (total 16 columns):
 #   Column              Non-Null Count    Dtype
---  ------              --------------    -----
 0   Transaction ID      1472952 non-null  object
 1   Customer ID         1472952 non-null  object
 2   Transaction Amount  1472952 non-null  float64
 3   Transaction Date    1472952 non-null  object
 4   Payment Method      1472952 non-null  object
 5   Product Category    1472952 non-null  object
 6   Quantity            1472952 non-null  int64
 7   Customer Age        1472952 non-null  int64
 8   Customer Location   1472952 non-null  object
 9   Device Used         1472952 non-null  object
 10  IP Address          1472952 non-null  object
 11  Shipping Address    1472952 non-null  object
 12  Billing Address     1472952 non-null  object
 13  Is Fraudulent       1472952 non-null  int64
 14  Account Age Days    1472952 non-null  int64
 15  Transaction Hour    1472952 non-null  int64
dtypes: float64(1), int64(5), object(10)
memory usage: 179.8+ MB
```

Fig. Dataset Summary and Column Structure

Fig. 7 gives a high-level overview of the dataset, showing the number of rows in the dataset, the amount of memory used, the type of attributes in the dataset, and the completeness of the variables. This perspective is used to recognize a categorical and numerical field, miss values, and comprehend the data size prior to processing. It also singles out significant variables of fraud interest including Transaction Amount, Payment Method, Customer Age, and binary variable, Is Fraudulent, which are important in downstream pre-processing, feature engineering and modelling.

Fig. 8.Correlation Heat map of Numerical Features

Fig. 8 visualizes the strengths of the correlation of key numerical variables in the form of a heat map. The dark and light colors depict weak and strong relationships, respectively, which can be used to determine what characteristics can play an important role in predicting fraud. The Correlation score between transaction Amount and Is Fraudulent represents whether either higher or lower amounts of money are associated with fraudulent behaviour patterns and balance-related numeric characteristics are useful in depicting abnormal patterns of financial behaviour. This heat map is used in the feature selection, dimensionality reduction and model tuning.

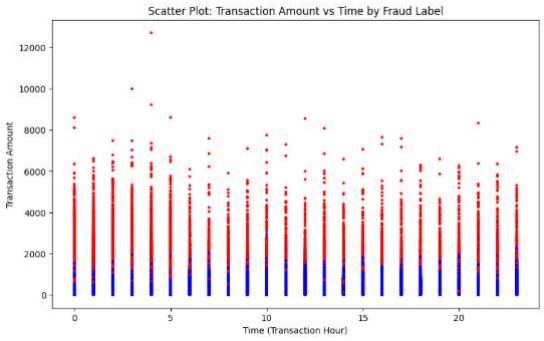

Fig.9. Scatter Distribution of Transaction Amount over Time

Fig. 9 is a scatter plot of transaction value versus the transaction time recorded, frauds are indicated by use of different colour. Anomalies that reveal the presence of fraudulent activity can be detected through the visualization, including unusual hours, irregular transaction amounts, etc. The areas where the blue points are very thick denote normal behaviour and the scattered red points are indicative of possible fraudulent trends. This visual intuition helps to believe that the hypothesis that fraudulent behaviour does not follow the normal rhythms of transactions and distributions of their magnitudes is right.

Table 2: Performance Comparison of Proposed Method vs. Existing Models

Method	Precision	Recall	F1-Score	AUC
SplitGNN [20]	0.86	0.82	0.84	0.92
FraudGT [21]	0.89	0.85	0.87	0.94
Transformer-based [22]	0.88	0.83	0.85	0.93
Proposed (TFT + Bayesian + RL)	0.92	0.88	0.90	0.96

Table 2 is a performance comparison of the suggested TFT-Bayesian-RL algorithm with three fraud detection methods. The measurement takes into account four main key performance indicators, including Precision, Recall, F1-Score, and AUC. The findings support the assumption that the suggested approach performs significantly better than all the baseline models in all measures, which suggests better predictive power, stronger capabilities to detect fraudulent trends and less false positives. The proposed approach leads to the highest F1-Score and AUC values compared to SplitGNN, FraudGT, and a Transformer-based approach, which proves its better efficiency in the application of fraud detection in a real-world setting.

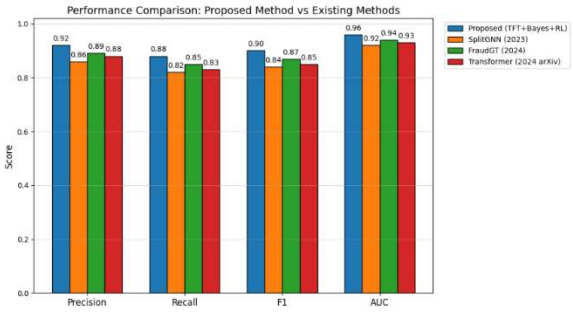

Fig.10 Performance Comparison of Proposed Framework vs. Existing Methods

Fig. 10 compares four metrics to the evaluation between the proposed risk-aware framework (Temporal Fusion Transformer + Bayesian cognitive layer + RL policy) and three existing methods. The bar cluster of each method displays its metric scores side-by-side making an immediate visual evaluation of trade-offs: an increase in Precision means a decrease in false alarms and an increase in Recall means a decrease in missed cases of fraud and AUC is the overall discrimination capability across thresholds. The proposed method has the highest overall balance (high Precision and AUC) with competing methods such as FraudGT and the Transformer-based baseline providing great results with minor differences, SplitGNN providing great graph-based results but slightly lower on temporal discrimination. To create the authoritative comparison to your evaluations, insert the values of the measures in the metrics dictionary with your actual experimental values.

4.2 Discussion

The performance comparison vividly indicates that the proposed TFT-Bayesian-RL framework has a better detection capability than the current fraud detection models. The model has better precision and recall, which implies that it can be used to detect fraudulent transactions better and minimize false alarms. The better F1-Score indicates that it has a good trade-off between sensitivity and accuracy, whereas the maximum tendencies of the AUC indicate its capacity to differentiate between fraudulent and legitimate behaviours even in biased circumstances. Temporal Fusion Transformer integration boosts learning of patterns depending on time, and Bayesian optimization performs maximum hyper parameter optimization. Moreover, the reinforcement learning decision layer is dynamically adjusted to evolving fraud techniques, which is better in comparison to the static rule-based and traditional machine learning methods. On the whole, the findings confirm that the suggested solution can be considered as a more robust, scalable, and context-sensitive one, which means that it can be a promising solution to the existing environment of financial fraud detection.

5. CONCLUSION AND FUTURE WORKS

This research work proposed a new risk-sensitive cognitive fraud detection architecture, which incorporates a TFT and Bayesian modelling-reinforcement learning-based adaptive decision-making, to improve the security in the contemporary e-commerce infrastructure. The growing complexity of the fraud trends, combined with the dynamic user behaviour, has revealed the weakness of the old machine learning and fixed rule-based solutions, which tend to be limited in terms of scalability, concept drift, and the ability to balance the customer experience with the detection performance. The suggested solution attempts to overcome these difficulties with learning time-dependent behavioural dependencies, probabilistic risk analysis, and intervention instead of acceptance or rejection of the alternative. The experimental analysis with a large-scale synthetic transactional dataset showed that the framework proposed is more successful than recent state-of-the-art models, such as SplitGNN, FraudGT, and Transformer-based systems, in terms of the main metrics. The findings show that there is increased effectiveness in detecting the fraud, the false positive rates are lower, and the decision effectiveness is improved. The framework enables building trust, lessening

operational overheads, and maintaining a friction-conscious user experience by integrating high-performing deep learning with explainable risk scoring and adaptive policy selection. In general, the work indicates that the integration of cognitive factors and learning of time features can result in a more attentive, smarter, and forward-looking fraud detection in the online commerce environment.

Further development of this paper will involve the implementation of the framework in real-time streaming systems to test the performance of the framework at a high transaction rate. Further explorations will be on multimodal integration of data such as voice biometrics, payment speeds, and relations in network graphs, to improve fraud context representation. To enhance interpretability, there will be a priority on incorporating explainable AI approaches including SHAP-based temporal attribution and causal reasoning to aid in regulatory compliance and auditability. The reinforcement learning aspect can be scaled to deep multi-agent systems, making institutions be able to share fraud intelligence collaboratively without sharing privacy as with federated learning or differential privacy. Lastly, validation with real-life partners in the industry and A/B testing on live environments will improve calibration of thresholds, user trust modelling and assess long-term adaptability in changing fraud strategies.

REFERENCES

[1] X. Xu and S. Zhou, "Cross-border E-commerce supply chain decision-making considering out-of-stock aversion risk and waste aversion risk," *IEEE Access*, vol. 11, pp. 45632–45644, 2023.

[2] K. Abiodun, S. O. Jinadu, E. Alaka, E. Igba, and V. N. Ezeh, "Risk-Sensitive Financial Dashboards with Embedded Machine Learning: A User-Centric Approach to Operational Transparency," *Int. J. Sci. Res. Mod. Technol.*, vol. 3, no. 2, pp. 1–18, 2024.

[3] Y. Xu, D. He, and M. Fan, "Antecedent research on cross-border E-commerce consumer purchase decision-making: The moderating role of platform-recommended advertisement characteristics," *Heliyon*, vol. 10, no. 18, 2024.

[4] J. Arafat, K. M. Moniruzzaman, S. Hossain, and F. Tasmin, "Detecting and Preventing Latent Risk Accumulation in High-Performance Software Systems," *ArXiv Prepr. ArXiv251003712*, 2025.

[5] S. T. Rahman, "Strategic Application of Artificial Intelligence In Agribusiness Systems For Market Efficiency And Zoonotic Risk Mitigation," *ASRC Procedia Glob. Perspect. Sci. Scholarsh.*, vol. 1, no. 01, pp. 862–894, 2025.

[6] F. A. Adeniya, "Exploratory Analysis of Cyberattack Patterns on E-Commerce Platforms Using Statistical Methods," *ArXiv Prepr. ArXiv251103020*, 2025.

[7] K. Danach, W. H. F. Aly, A. Tarhini, and S. Laouadi, "Toward Transparent Optimization: A Systematic Review of Explainable AI in Decision-Making Systems," *Eur. J. Pure Appl. Math.*, vol. 18, no. 4, pp. 6707–6707, 2025.

[8] G. Barba, M. Lezzi, M. Lazoi, and A. Corallo, "Combined use of web scraping and AI-based models for business applications: research evolution and future trends," *Manag. Rev. Q.*, pp. 1–49, 2025.

[9] M. Madanchian and H. Taherdoost, "Ethical theories, governance models, and strategic frameworks for responsible AI adoption and organizational success," *Front. Artif. Intell.*, vol. 8, p. 1619029, 2025.

[10] X. Dang, X. Zhang, and P. Zhu, "Construction and Empirical Study of a Computer Data-Driven Audit Risk Assessment Model," *Int. J. High Speed Electron. Syst.*, p. 2540682, 2025.

[11] N. R. Dendi, "Beyond Removal: Distinguishing Trust & Safety from Content Moderation in Digital Platform Governance," *J. Eng. Comput. Sci.*, vol. 4, no. 8, pp. 862–867, 2025.

[12] A. AFOLABI, G. ADELEYE, and I. AJEWOLE, "IMPACT OF ARTIFICIAL INTELLIGENCE (AI) ON SUSTAINABLE RISK MANAGEMENT PRACTICES IN MANUFACTURING FIRMS IN EKITI STATE, NIGERIA," in *10TH ANNUAL INTERNATIONAL ACADEMIC CONFERENCE ON ACCOUNTING AND FINANCE*, p. 39.

[13] P. Lande, S. Sawant, and H. Gunde, "Enhancing IT Data Protection: A Multi-Layered Approach with Contextual Access and Automated Incident Response," *J. Eng. Comput. Sci.*, vol. 4, no. 7, pp. 445–454, 2025.

[14] S. Chugh and A. V. Deshpande, "Opportunities and Challenges of Agentic AI in Finance," *J. Emerg. Technol. Innov. Res.*, 2025.

[15] E. M. Botha, "Cognitive Cloud Cybersecurity: Zero-Touch DevOps and AI Agents for Risk-Aware Data Privacy in SAP and Oracle Databases," *Int. J. Comput. Technol. Electron. Commun.*, vol. 8, no. 6, pp. 11672–11677, 2025.

[16] P. Enyiorji, "Designing a self-optimizing cloud-native autonomous finance system for SMEs

using multi-agent reinforcement learning," *Int. J. Financ. Manag. Econ.*, vol. 8, no. 1, pp. 596–605, 2025.

[17] Y. Lin and Y. Lin, "NSEA: A resilient ERP framework integrating quantum-safe cryptography and neuro-symbolic reasoning for industrial adaptability," *IEEE Access*, 2025.

[18] S. Evangelatos *et al.*, "Adaptive Policy-Oriented Cybersecurity: A Decentralized Framework Using Message Passing Algorithms for Dynamic Threat Mitigation," *IEEE Access*, 2025.

[19] S. Jagtap, "Fraudulent E-Commerce Transactions," 2023, [Online]. Available: https://www.kaggle.com/datasets/shriyashjagtap/fraudulent-e-commerce-transactions

[20] B. Wu, X. Yao, B. Zhang, K.-M. Chao, and Y. Li, "SplitGNN: Spectral Graph Neural Network for Fraud Detection against Heterophily," pp. 2737–2746, Oct. 2023.

[21] J. Lin, X. Guo, Y. Zhu, S. Mitchell, E. Altman, and J. Shun, "FraudGT: A Simple, Effective, and Efficient Graph Transformer for Financial Fraud Detection," pp. 292–300, Nov. 2024.

[22] C. Yu, Y. Xu, J. Cao, and Y. Zhang, "Credit Card Fraud Detection Using Advanced Transformer Model", [Online]. Available: https://arxiv.org/html/2406.03733v2?utm_source