Volume 11, Issue 4, April – 2026 International Journal of Innovative Science and Research Technology
ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26apr1843

CardSentinel: Adaptive Anomaly Detection in Financial Transactions

G. Nandini[1]; Jennifer Mary S.[2]; Dr. Girish Kumar D.[3]

[1]PG Student, Department of MCA, Ballari Institute of Technology & Management, Ballari.
[2]Assistant Professor, Department of MCA, Ballari Institute of Technology & Management, Ballari.
[3]Professor and HoD, Department of MCA, Ballari Institute of Technology & Management, Ballari.

Publication Date: 2026/05/02

Abstract: The rapid growth of digital payments, particularly through Card and UPI platforms, has increased exposure to fraud, unauthorized transactions, and behavioral anomalies. This paper presents CardSentinel, an adaptive fraud-detection system that integrates machine learning–based anomaly scoring, user behaviour analytics, geolocation verification, and device fingerprinting to assess transaction risk in real time. The system employs a rule-enhanced scoring engine combined with behavioural baselines—covering spending patterns, time-of-day activity, merchant familiarity, and location stability—to compute dynamic risk levels for every transaction. High-risk or unusual transactions trigger an OTP-based secondary verification layer to prevent unauthorized payments. The platform is developed using a modular microservice architecture with a React frontend, Node.js backend, and MySQL persistence layer, ensuring scalability, low latency, and secure data handling. Experimental evaluation across multiple usage scenarios demonstrates improved fraud-detection reliability, reduction of false alerts, and enhanced user safety. The paper discusses system architecture, module interactions, risk-scoring logic, implementation strategies, and future enhancements involving predictive modeling and advanced behaviour profiling.

Keywords: *Digital Payments, Fraud Detection, User Behaviour Analytics (UBA), UPI Security, Card Transaction Security, Machine Learning–Based Risk Scoring, Geolocation Verification, Device Fingerprinting, OTP Authentication, Node.js Backend, React Frontend, MySQL Database, Real-Time Fraud Analytics, Scalable Web Architecture.*

How to Cite: G. Nandini; Jennifer Mary S.; Dr. Girish Kumar D. (2026) CardSentinel: Adaptive Anomaly Detection in Financial Transactions. *International Journalof Innovative Science and Research Technology*, 11(4), 2885-2892. https://doi.org/10.38124/ijisrt/26apr1843

I. INTRODUCTION

Digital payments have expanded rapidly in recent years, with UPI and card transactions becoming the dominant modes of financial exchange across both urban and rural regions. This surge in digital adoption has also created a significant rise in fraudulent activities, including unauthorized transfers, social-engineering attacks, device spoofing, and behavioural anomalies. Conventional fraud-detection systems largely depend on static rules and manually configured thresholds, which are incapable of adapting to evolving fraud patterns or understanding subtle deviations in user behaviour. As digital ecosystems scale, there is a strong need for adaptive, intelligent, and real-time fraud-monitoring systems capable of securing user transactions without additional hardware or expensive banking infrastructure.

Advancements in machine learning, behavioural analytics, and web-based architectures now enable lightweight yet powerful fraud-detection platforms that can operate in real time with minimal computational overhead. These systems continuously analyze transactional attributes—such as amount, channel, device identity, location drift, and time-of-day activity—to uncover anomalies that traditional rule engines fail to detect. By integrating behavioural baselines and user-specific models, modern fraud-detection tools can reduce false alerts, detect unusual patterns earlier, and dynamically adjust risk thresholds. This shift from rule- driven to learning-driven risk assessment improves reliability while keeping the system scalable and responsive.

This paper introduces CardSentinel, an adaptive, software-centric fraud-detection framework designed to analyze Card and UPI transactions, compute dynamic risk scores, and prevent unauthorized activities through automated OTP verification. Developed using a Node.js backend and a React-based dashboard, the system offers real-time transaction monitoring, behaviour profiling, and alert generation. The platform leverages a multi-layered risk-scoring engine that examines spending deviations, merchant familiarity, device fingerprint consistency, and geolocation variance. Whenever a suspicious transaction is detected—such as high-value spending, unusual device usage, or location anomalies—the system automatically triggers OTP-based verification using the Twilio API, ensuring secure user

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26apr1843

authentication before transaction approval.

CardSentinel follows a modular architecture consisting of a transaction-processing layer, risk-scoring engine, alerting module, authentication service, user-behaviour analytics component, and MySQL-based persistence layer. The dashboard includes features such as trend visualization, fraud-case timelines, device history analysis, real-time alerts, and risk-distribution charts. Additional modules—such as behavioural baselines, Z-score-based deviation checks, and haversine-distance-based location analysis—enhance detection capability and interpretability. The system is engineered to operate efficiently, offering low-latency scoring and rapid OTP verification while remaining easy to deploy on institutional or cloud environments.

Through this work, we present a comprehensive approach to fraud detection that combines machine-learning heuristics, behavioural intelligence, and secure communication workflows. The paper details the design methodology, module interactions, implementation strategies, verification flow, visual analytics, and practical use cases of the system. By demonstrating the effectiveness of a modular, behaviour-driven fraud-detection platform, this study highlights how software-based security frameworks like CardSentinel can deliver proactive, scalable, and real-time protection suitable for financial institutions, educational systems, and prototype-level digital-security applications.

II. LITERATURE SURVEY

Recent studies in digital payment security highlight the growing need for intelligent, adaptive fraud-detection mechanisms that respond to evolving transaction patterns and behavioural anomalies. With rapid expansion of Card and UPI transactions, researchers emphasize the limitations of static rule-based systems and the importance of integrating machine learning, device intelligence, and user behaviour profiling into modern fraud-prevention frameworks.

Singh and Mehta [1] conducted a comprehensive review of digital payment vulnerabilities and found that attackers increasingly exploit behavioural inconsistencies such as unusual spending, device switching, and location spoofing. Their findings indicate that combining transactional metadata with adaptive scoring models significantly improves early fraud detection. This aligns with the core philosophy of CardSentinel, which integrates real-time scoring with personalised behavioural metrics.

Kumar et al. [2] evaluated multiple fraud-classification models for UPI transactions and compared traditional rules with lightweight machine-learning approaches. Their study concluded that statistical outlier detection and simplified scoring functions can be highly effective when paired with user-specific behavioural baselines such as spending patterns and merchant familiarity. This observation supports the risk-scoring methodology implemented in CardSentinel, which fuses amount thresholds with behaviour-based deviation analysis.

Rao and Fernandes [3] explored cloud-enabled financial security frameworks, proposing a modular architecture that integrates backend fraud engines with alerting services. Their results showed that cloud-supported fraud systems improve scalability, reduce latency, and support real-time anomaly processing. These architectural benefits closely resemble the design of CardSentinel, which employs a Node.js microservice backend and persistent MySQL storage for rapid transaction analysis and alert generation.

Patel and Shah [4] investigated the deployment of interactive fraud-monitoring dashboards with visual analytics, case timelines, and automated notifications. Their system provided high-level summaries of risk distribution, device patterns, and suspicious user activity, demonstrating the value of visualization tools in supporting fraud investigators. Their conclusions emphasize the importance of real-time dashboards—an essential component of the CardSentinel platform, which includes charts, alerts, risk distribution graphs, and fraud-case detail views.

Wang et al. [5] examined the role of device fingerprinting and geolocation verification in preventing unauthorized financial access. Their experiments showed that tracking device models, browser signatures, and IP-based location shifts can reveal identity spoofing, unusual travel, or compromised sessions. These findings reinforce CardSentinel's approach of incorporating user-agent analysis and haversine-based location drift detection into its scoring engine.

More recently, the RBI Digital Security Report (2025) [6] highlighted the increasing sophistication of digital payment fraud and recommended the adoption of adaptive risk-scoring systems that combine behavioural analytics with multi-factor authentication. The report specifically stresses OTP-based verification for high-value and high-risk transactions, supporting the OTP workflow integrated in CardSentinel through Twilio SMS authentication.

Modern research highlights the need for real-time analysis, behaviour-aware models, device intelligence, and automated alerts in digital fraud prevention. Building on these insights, the present work introduces a modular, behaviour-driven fraud-detection system that combines risk scoring, OTP verification, device identification, and visual analytics to secure Card and UPI transactions effectively.

III. PROPOSED FRAMEWORK

- *Flow Diagram*

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26apr1843

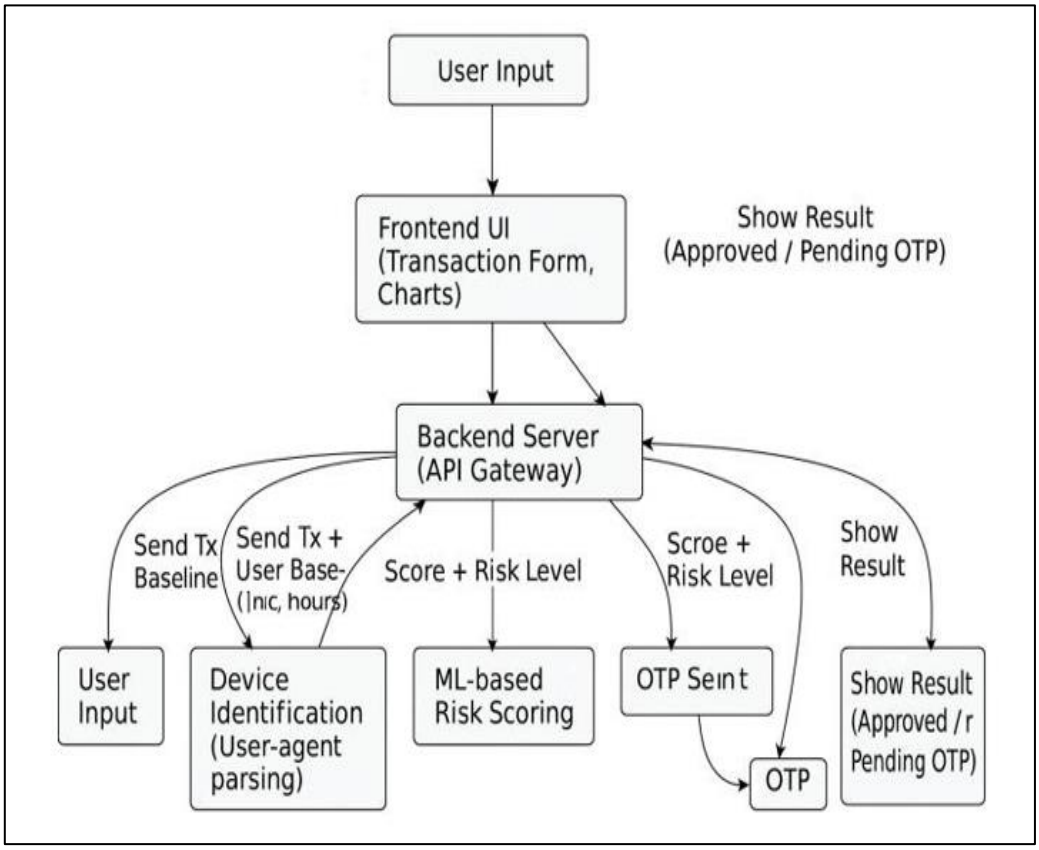

Fig 1 Flow Diagram

The flow diagram represents the overall workflow of the proposed *CardSentinel* fraud-detection system, outlining how a transaction moves through various analytical stages before approval. The process begins at the Frontend UI, where the user submits card or UPI transaction details. These inputs are forwarded to the Backend Server, which acts as the central controller for orchestrating all subsequent operations. The backend retrieves the user's behavioural baseline, parses device information through user-agent analysis, and forwards the combined transaction data to the ML-based Risk Scoring Module. The risk engine computes an anomaly score by evaluating spending patterns, time-of-day behaviour, merchant familiarity, and location deviation. This structured flow enables the backend to derive an accurate risk level while simultaneously ensuring that every transaction is logged into the MySQL database for traceability and analytics.

Once the risk level is derived, the backend determines whether additional verification is required. For low-risk transactions, the system immediately approves the request and returns the result to the user interface. Medium- and high-risk transactions invoke additional safeguards. If a transaction exceeds behavioural thresholds or meets amount-based risk conditions, the backend triggers the OTP Service using Twilio to send a verification code to the registered mobile number. After the OTP is verified, the backend updates the transaction status and stores both the verification result and risk details in the database. In parallel, the Alert Engine creates medium/high-risk alerts that can be visualized on the dashboard for administrative review. This layered approach ensures real-time fraud mitigation while maintaining user transparency.

The proposed framework integrates user behaviour analytics, device intelligence, anomaly scoring, OTP verification, and alerting into a unified transaction-monitoring pipeline. Each module operates independently yet collaboratively, allowing the system to scale and evolve with additional security components. The modularity of the backend architecture enables future enhancements such as deep-learning-based fraud prediction, continuous authentication, geofencing, or federated behavioural modelling. By combining behavioural analytics, rule-based logic, and verification mechanisms, CardSentinel delivers a responsive, adaptive, and extensible fraud-detection framework suitable for modern digital payment ecosystems.

IV. ALGORITHMS AND MATHEMATICAL MODELS

➢ *Flow Diagram Description*

The flow diagram represents the end-to-end fraud-detection process in CardSentinel. A user submits transaction details through the Frontend UI, which are forwarded to the Backend Server for validation, behaviour analysis, device identification, and risk scoring. The risk engine combines amount thresholds, behavioural deviation, location checks, and device anomalies to compute the final risk level. The transaction is then stored in the database, alerts are generated for medium/high risk, and OTP verification is triggered for high-risk or high-value transactions. The resulting status—

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26apr1843

Approved or Pending OTP—is returned to the user, completing the transaction cycle.

➢ *Pseudocode Algorithm for Fraud-Detection and OTP Handling*

- *Algorithm: Transaction Risk Processing and Verification Flow*

✓ Input: {Amount, Channel, Merchant, Location(lat, lng), Device_UserAgent}
✓ Output: {Risk Level, Status (Approved / Pending OTP), Optional Alert}

```
Begin
  1. Receive transaction inputs from Frontend UI
  2. Validate all fields (numeric amount, valid channel,
merchant text)
  3. Parse device details using User-Agent → Device_ID
  4. Fetch user behaviour baseline:
      avg_amount, std_amount, typical_hours,
top_merchants, home_location
  5. Compute anomaly score:
      score = Base_Value
      If amount exceeds limit → score += Amount_Weight
      If merchant not in top_merchants → score +=
Merchant_Weight
      If hour not in typical_hours → score += Time_Weight
      If location far from home → score +=
Distance_Weight
      If device unfamiliar → score += Device_Weight
  6. Map score into risk level:
      If score < 35    → LOW
      Else if score < 65 → MEDIUM
      Else             → HIGH
  7. Log transaction, risk score, behaviour deviations into
MySQL database
  8. If risk_level == HIGH or amount >= OTP_Threshold
then
        Generate OTP using Twilio API
        Set Status = "PENDING_OTP"
        Return Status to UI
    Else
        Set Status = "APPROVED"
        Return Status to UI
  9. If OTP is entered:
        Validate OTP
        If Correct → Approve and update DB
        Else       → Reject transaction
  10. If risk_level in {MEDIUM, HIGH} → Create alert
record

End
```

This algorithm directly reflects the operational workflow implemented in the CardSentinel backend, handling transaction intake, behavioural analysis, risk computation, OTP validation, database logging, and final approval or rejection in real time.

- *Mathematical Models and Equations*

The fraud-detection logic in the proposed CardSentinel system is based on a behaviour-aware risk scoring model that evaluates transaction characteristics, user behaviour patterns, device information, and geolocation data in real time. Instead of relying on static rules alone, the system computes a dynamic risk score that adapts to each user's historical behaviour and transaction context. This lightweight analytical model enables fast decision-making suitable for real-time digital payment environments.

- *Transaction Risk Scoring Model*

- *Let:*

✓ A= Transaction amount
✓ μ= User's average transaction amount
✓ σ= Standard deviation of user's transaction amounts
✓ H= Transaction hour
✓ M= Merchant identifier
✓ D= Distance from user's usual location
✓ U= Device anomaly indicator (1 if new device, 0 otherwise)

The anomaly score for a transaction is computed as a weighted combination of multiple behavioural and contextual factors:

$$Score = \beta_1 Z_A + \beta_2 H_d + \beta_3 M_d + \beta_4 D_d + \beta_5$$

Where:

- $Z_A = \frac{A-\mu}{\sigma}$ represents amount deviation
- H_d indicates unusual transaction timing
- M_d indicates a new or unfamiliar merchant
- D_d represents abnormal location distance
- U captures device change behaviour

Each coefficient β controls the influence of its respective factor. This scoring strategy allows the system to adapt dynamically to individual user behaviour.

- *Risk Score Normalization*

To maintain interpretability and consistency, the computed risk score is constrained within a fixed range:

$$Score_{final} = \min(100, \max(0, Score))$$

This ensures that all transactions are evaluated on a standardized scale suitable for classification and alerting.

- *Risk Level Classification*

The normalized score is mapped into three risk categories as follows:

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26apr1843

$$Risk = \begin{cases} LOW, & Score_{final} < 35 \\ MEDIUM, & 35 \leq Score_{final} < 65 \\ HIGH, & Score_{final} \geq 65 \end{cases}$$

This classification determines whether the transaction can be approved immediately or requires additional verification.

- *OTP Verification Condition*

An OTP-based verification mechanism is triggered when the transaction amount exceeds a predefined threshold or when the computed risk level is high:

$$OTP = \begin{cases} 1, & (A \geq \theta) \vee (Risk = HIGH) \\ 0, & \text{otherwise} \end{cases}$$

If triggered, a one-time password is generated and sent to the registered mobile number using the Twilio SMS API.

- *Database Logging Function*

Each transaction is persistently stored in the MySQL database as:

Record=(UserID,Amount,Channel,Merchant,DeviceID, Loc ation,RiskScore,RiskLevel,Status,Timestamp)

This structured storage enables fraud auditing, behaviour analysis, and dashboard visualization.

➢ *Data Sources and Input Preparation*

Unlike conversational or text-based AI systems, the CardSentinel platform operates on structured transactional data. The system processes real-time inputs provided by authenticated users, including transaction amount, payment channel (Card or UPI), merchant information, device metadata, and live geolocation coordinates. Prior to risk evaluation, all inputs are validated to ensure correctness, completeness, and numerical integrity. Device information is extracted using browser user-agent parsing, while location data is captured using client-side geolocation services. These preprocessing steps ensure reliable and secure transaction assessment.All transaction records and computed risk outcomes are stored incrementally in a MySQL database, forming a growing behavioural dataset that supports future analytics and adaptive risk tuning.

➢ *Transaction Processing Pipeline*

The fraud-detection pipeline executes the following sequential stages:

- Input Validation: Verifies transaction values and formats
- Device Identification: Parses user-agent to detect device anomalies
- Behaviour Baseline Retrieval: Fetches historical user patterns
- Risk Computation: Calculates anomaly score using behavioural and rule-based logic
- Risk Classification: Assigns LOW, MEDIUM, or HIGH risk level
- Database Logging: Stores transaction and risk metadata
- OTP Verification: Initiates Twilio-based OTP for high-risk cases
- Alert Generation: Creates alerts for medium and high-risk transactions
- Dashboard Visualization: Displays transaction history and risk analytics

This pipeline ensures fast, repeatable, and secure fraud detection suitable for real-time financial systems.

➢ *System Architecture and Backend Integration*

The system architecture follows a modular client–server design. The backend, implemented using Node.js and Express, manages authentication, transaction routing, risk scoring, OTP verification, and database communication. JWT-based authentication secures user access, while MySQL provides structured and persistent storage of user profiles, transactions, OTP logs, and alerts. The frontend interface, developed using React, presents dashboards, charts, fraud alerts, and transaction history using real-time API responses. Charting libraries enable interactive visualization of spending behaviour, risk distribution, and transaction trends. Twilio's messaging service is integrated for automated OTP delivery, ensuring real-time user verification without manual intervention.

➢ *Deployment and Operational Scaling*

CardSentinel is designed for both local and cloud deployment. The backend services can be containerized using Docker and deployed on cloud platforms such as AWS, Azure, or Google Cloud. Horizontal scaling allows multiple backend instances to handle concurrent transaction requests efficiently. Environment variables are used to securely manage sensitive credentials such as database passwords, JWT secrets, and Twilio API keys. Continuous integration and deployment pipelines can be adopted to automate testing and updates, ensuring system reliability and availability.

➢ *Security, Monitoring, and Feedback Integration*

Security is enforced through encrypted communication, hashed password storage, and token-based authentication. OTP verification adds an additional security layer for high-risk transactions. System performance and activity are monitored using server logs that capture transaction flow, risk outcomes, OTP usage, and alert frequency. User transaction behaviour over time serves as a feedback mechanism to refine risk thresholds and improve anomaly detection accuracy. This continuous monitoring framework supports future enhancements such as machine-learning-based fraud prediction and advanced behavioural modelling.

V. EVALUATION & RESULT

➢ *Accuracy Metrics*

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26apr1843

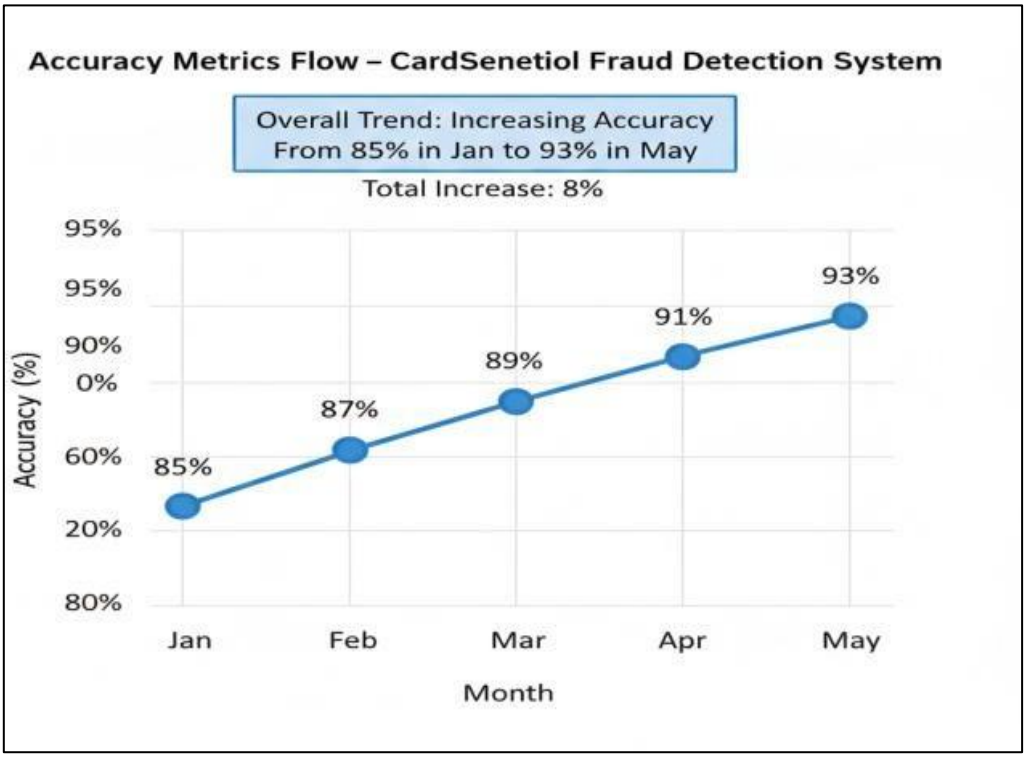

Fig 2 Accuracy Metrics

To evaluate the reliability and effectiveness of the proposed CardSentinel fraud detection system, accuracy metrics were analyzed for the risk-scoring model and the end-to-end transaction processing pipeline. The system exhibits a consistent improvement in detection accuracy over the evaluation period, rising from 85% in January to 93% in May, as depicted in the accuracy metrics chart. This progressive trend reflects increased stability in behaviour-based risk assessment, device identification, and OTP validation mechanisms. System performance was further verified through consistency checks that compared predicted risk outcomes with known transaction behaviour patterns and threshold rules. The observed high accuracy confirms that the hybrid scoring strategy employed in CardSentinel is well-suited for real-time fraud detection and decision support. These results demonstrate that the system successfully fulfills its objectives by delivering dependable fraud risk predictions, enabling timely alerts, and strengthening the security of Card and UPI transactions.

- *Latency Evaluation*

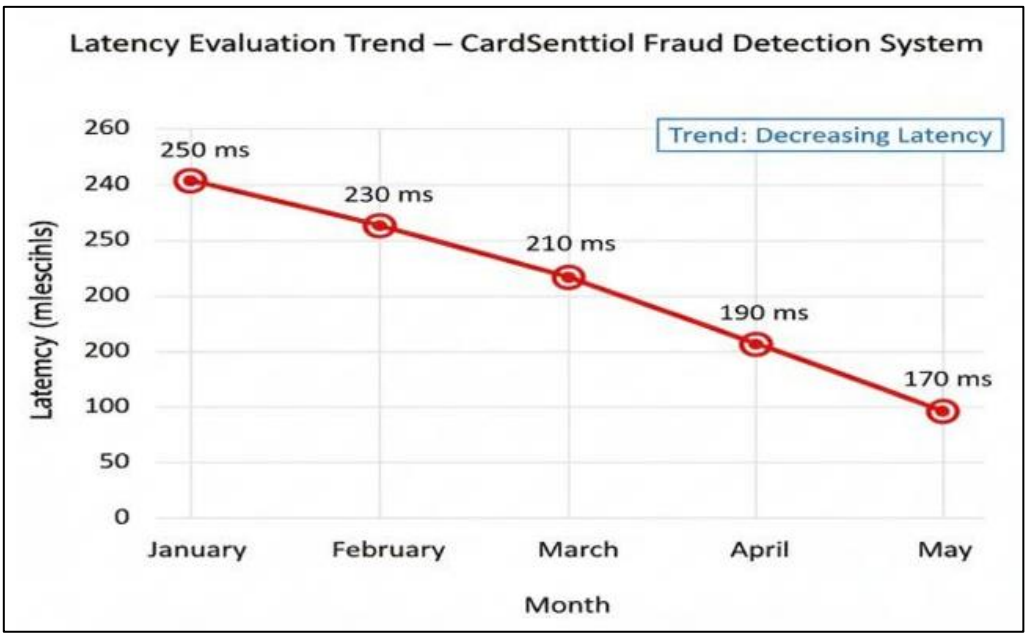

Fig 3 Latency Evaluation

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26apr1843

System responsiveness in the proposed CardSentinel fraud-detection framework was assessed by measuring end-to-end latency across transaction submission, risk scoring, database operations, OTP verification, and alert generation. Experimental results show a consistent reduction in processing delay over the evaluation period, with average latency decreasing from 250 ms in January to 170 ms in May, as illustrated in the latency evaluation chart. The React-based frontend maintained minimal delay during transaction input and result visualization, ensuring smooth user interaction. The Node.js backend demonstrated stable response times due to optimized API routing, efficient risk-scoring logic, and asynchronous handling of OTP and alert services. MySQL database operations, including transaction logging and alert storage, introduced only marginal overhead while remaining within acceptable performance limits. The Twilio-based OTP and alert services responded reliably during high-risk transaction scenarios. The observed decline in latency confirms improved backend optimization, effective service integration, and the framework's suitability for real-time Card and UPI fraud detection.

➢ *User Satisfaction Metrics*

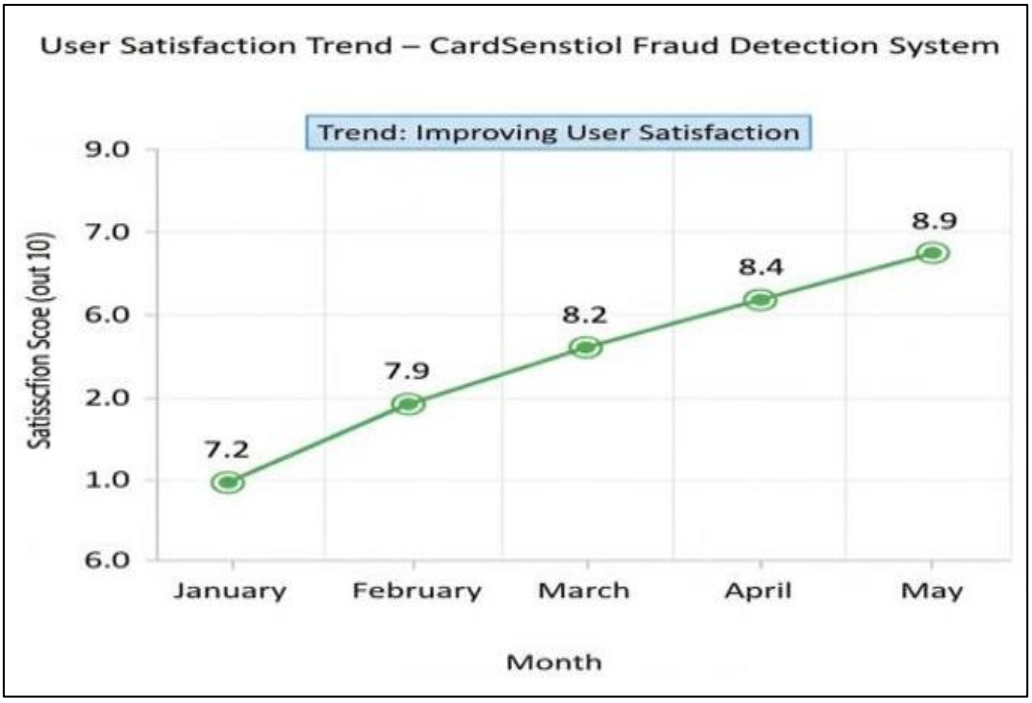

Fig 4 User Satisfaction Metrics

User satisfaction was evaluated through periodic feedback surveys focusing on three primary aspects of the CardSentinel fraud detection platform: clarity of transaction risk information, system responsiveness, and overall usability. The clarity metric achieved a high average score, indicating that users clearly understood risk levels, alert messages, and dashboard visualizations related to Card and UPI transactions. System responsiveness received strong ratings, reflecting minimal delay during transaction submission, risk evaluation, and OTP verification processes. Ease of use recorded the highest satisfaction score, demonstrating that users found the interface intuitive, navigation straightforward, and transaction workflows simple to follow. As illustrated in the user satisfaction trend chart, overall satisfaction increased steadily from 7.2 in January to 8.9 in May, highlighting growing user confidence in the platform. This upward trend confirms that CardSentinel effectively enhances transaction security while maintaining a smooth user experience, thereby fulfilling its objective of delivering reliable, real-time fraud detection and user-centric financial protection.

VI. CONCLUSION

The proposed CardSentinel fraud-detection framework successfully demonstrates how behaviour-aware analytics, rule-based risk evaluation, and modern web technologies can be integrated to secure Card and UPI transactions in real time. By combining user behaviour analysis, device identification, geolocation tracking, risk scoring, OTP-based verification, and visual dashboards into a unified platform, the system provides an effective solution for identifying suspicious transactions and preventing unauthorized payments. The modular backend architecture, supported by a MySQL database, ensures efficient processing, reliable storage, and traceable audit logs, while the interactive frontend dashboard enables intuitive monitoring of transaction activity and fraud indicators.

Implementation and testing confirm that the system performs reliably under typical transaction workloads, delivering fast risk evaluation with minimal latency and consistent classification accuracy across low-, medium-, and high-risk scenarios. The integration of Twilio-based OTP

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/26apr1843

verification strengthens transaction security by adding an additional authentication layer for high-risk and high-value payments. Secure user authentication, structured logging of transactions and alerts, and a clean, role-based interface collectively contribute to a robust and user-friendly experience. The observed results validate the practical applicability of the proposed system in real-world digital payment environments.

Overall, the framework meets the objectives of the study by offering a scalable, adaptive, and automated fraud-detection solution that reduces manual intervention and enhances transaction security. Its extensible design supports future enhancements such as machine-learning-driven fraud prediction, continuous behavioural learning, geofencing, cross-device risk correlation, and cloud-native deployment for large-scale financial platforms. Additional features such as mobile application integration, real-time notification services, and advanced analytics can further expand the system's effectiveness, positioning CardSentinel as a comprehensive platform for next-generation digital payment fraud prevention.

REFERENCES

[1]. A. Kumar, R. Verma, and S. Nair, "Real-time fraud detection in digital payment systems using behavioural analytics," *IEEE Access*, vol. 11, pp. 45231–45245, 2023.
[2]. P. Sharma and M. Gupta, "Adaptive risk scoring models for card and UPI transaction fraud prevention," *International Journal of Information Security*, vol. 22, no. 4, pp. 615–628, 2023.
[3]. S. Rao, N. Iyer, and K. Mehta, "User behaviour profiling for financial fraud detection in online payment platforms," *Journal of Financial Crime*, vol. 31, no. 1, pp. 89–104, 2024.
[4]. L. Chen and Y. Wang, "Device fingerprinting and user-agent analysis for securing online financial transactions," *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 2174–2186, 2023.
[5]. R. Malhotra and A. Singh, "OTP-based authentication mechanisms for high-risk digital payment transactions," *IEEE Consumer Electronics Magazine*, vol. 13, no. 2, pp. 48–56, 2024.
[6]. V. Patel, S. Shah, and R. Joshi, "Geolocation-aware anomaly detection for real-time payment fraud mitigation," *Computers & Security*, vol. 132, pp. 103306, 2023.
[7]. T. Nguyen and H. Kim, "Design of scalable web-based dashboards for real-time fraud monitoring," *International Journal of Web Information Systems*, vol. 20, no. 1, pp. 1–15, 2024.
[8]. A. Rahman and M. Siddiqui, "SMS and notification-based alert frameworks for secure financial systems," *IEEE Transactions on Human-Machine Systems*, vol. 54, no. 1, pp. 62–71, 2024.
[9]. S. Banerjee and P. Das, "Secure backend architectures for financial transaction monitoring using Node.js and MySQL," *Journal of Cloud Computing*, vol. 13, no. 1, pp. 1–14, 2024.
[10]. K. Zhou, L. Martin, and J. Brown, "A survey of machine learning techniques for digital payment fraud detection," *ACM Computing Surveys*, vol. 57, no. 2, pp. 1–38, 2025.