The Journal of Computer Science and Its Applications. An International Journal of the Nigeria Computer Society (NCS)
Vol.32, No.1, June, 2026

Real-Time AI-Driven Behavioral Modeling in 3-D Secure ACS for Bot Fraud Detection

T. Eskala*
vikram9123@gmail.com

Engineering Services Architect
Mountain House, CA 95131, USA

Abstract
Automated checkout bots powered by generative AI and headless browsers now account for 47% of global e-commerce traffic, resulting in annual fraud losses exceeding $41.2 billion for card issuers worldwide. Traditional defenses such as CAPTCHAs, device fingerprinting, and static rules have become largely ineffective against modern adversarial bots that use reinforcement learning and diffusion models to generate human-like behavioral sequences. This paper presents a production-scale, continuously learning AI behavioral modeling framework deeply embedded within the Access Control Server (ACS)—the mandatory authentication gateway in the EMV 3-D Secure 2.3 protocol. By capturing and analyzing 42 high-resolution micro-behavioral signals including keystroke dynamics, pointer kinematics, touch pressure variance, hesitation entropy, and focus/blur patterns, a real-time ensemble combining LightGBM, 1D convolutional neural networks, and third-party behavioral biometric engines (BioCatch, BehavioSec) achieves bot detection with ROC AUC 0.96 and p95 latency of 42 ms. Live A/B deployment across seven Tier-1 European card issuers processing over 180 million transactions per month from January to June 2025 demonstrated a statistically significant 15% reduction in fraud incidents ($p < 0.001$), a 10% increase in legitimate transaction approval rates, an 8% decline in checkout abandonment, and annual cost savings exceeding $2.1 million—all while maintaining full compliance with PSD2 Strong Customer Authentication (SCA) requirements through formal classification of the learned behavioral profile as an inherence factor under RTS Article 18.

Keywords: AI Behavioral Modeling, Keystroke Dynamics, Ensemble Learning, Real-Time Fraud Detection, Adversarial Robustness, Continuous Learning, 3-D Secure, Bot Mitigation

1 Introduction

The rapid evolution of automated checkout bots has transformed payment fraud from a rule-based nuisance into a sophisticated AI-driven arms race. Modern bots leverage headless Chrome instances controlled by Puppeteer or Playwright, inject synthetic input events with sub-millisecond precision, randomize timing jitter using Gaussian distributions, and even employ generative adversarial networks (GANs) trained on real user sessions to produce indistinguishable behavioral sequences. Industry reports indicate that bot-initiated transactions have 3.2× higher average order value and account for nearly half of all e-commerce traffic in 2024–2025 (Imperva 2024; LexisNexis Risk Solutions 2024).

The 3-D Secure 2.3 protocol, governed by EMVCo, mandates that all authentication decisions route through the issuer's Access Control Server (ACS) before final card authorization. This strategic position makes the ACS the ideal location for injecting advanced behavioral intelligence. This paper introduces an AI-native behavioral modeling layer that transforms the ACS from a reactive checkpoint into a predictive, continuously learning fraud engine—without introducing user friction or violating PSD2 regulatory constraints.

The significance of this approach lies in its passive nature: unlike explicit challenges (OTP, CAPTCHA)

that increase abandonment by 8–12%, our system operates silently in the background, improving both security and user experience simultaneously. The framework is designed for production scale, handling 180 million transactions per month with minimal latency impact.

2 Threat Model: Bot Sophistication Levels

Contemporary checkout bots operate at multiple sophistication levels, each requiring increasingly advanced detection mechanisms:

- **Level 1: Simple HTTP-based scripts** — Easily blocked by rate limiting and basic device checks.
- **Level 2: Headless browsers with randomized delays and user-agent rotation** — Use Puppeteer/Playwright to simulate browser environments and evade fingerprinting.
- **Level 3: ML-enhanced bots using behavioral mimicry trained on real user data** — Incorporate reinforcement learning to generate natural timing variations and hesitation patterns.
- **Level 4: Adversarial bots specifically designed to poison or evade behavioral detection systems** — Employ GANs and diffusion models to synthesize sequences that actively target known detection features.

Our system was designed and validated against Level 3 and Level 4 adversaries, including GAN-generated input sequences and emulator farms operating tens of thousands of virtual devices simultaneously. Adversarial attacks include timing jitter injection, synthetic hesitation patterns, and coordinated poisoning attempts. Our defense strategy incorporates robust anomaly detection mechanisms such as behavioral entropy thresholds and cross-signal consistency checks, maintaining detection accuracy above 94% even under heavy simulated attacks.

3 High-Resolution Behavioral Data Acquisition

A lightweight (12 KB gzipped) JavaScript agent is injected into issuer-hosted 3-D Secure challenge pages and merchant checkout templates. The agent captures raw input events at 50 Hz, including:

- Keystroke press/release timestamps
- Pointer movement coordinates, velocity, acceleration and curvature
- Touch pressure and swipe vectors (mobile)
- Focus/blur and tab visibility state changes
- High-resolution timing via Performance.now()

All events are streamed in real time to a Kafka topic (acs.events.raw) with exactly-once semantics. Batching occurs every 100 ms to minimize network overhead while preserving real-time fidelity. This high-frequency sampling is critical for capturing subtle human motor control nuances (e.g., micro-hesitations, natural acceleration curves) that scripted bots struggle to replicate authentically.

The agent is designed for zero performance impact: asynchronous event buffering, no blocking calls, and automatic fallback to basic device data if injection fails.

4 System Architecture

4.1 Feature Engineering: 42 Behavioral Features Extracted

From raw event streams, 42 behavioral features are extracted per session, providing a high-dimensional representation of user intent. These features span four primary categories:

1) **Keystroke Dynamics (15 features):** Dwell time (duration key held), flight time (interval between keys), n-graph latencies (di/tri-graph timing), typing rhythm entropy, backspace ratio, capitalization patterns, and statistical moments (mean, variance, skewness, kurtosis) of timing distributions.
2) **Pointer Kinematics (12 features):** Velocity profile (mean, variance), acceleration peaks, curvature distribution, jerk magnitude, directional changes, and smoothness metrics (e.g., log-dimensionless jerk).
3) **Touch Gestures (8 features):** Pressure variance across taps, swipe velocity, swipe curvature, pinch/zoom patterns, and multi-touch consistency.
4) **Hesitation and Session Patterns (7 features):** Idle burst duration and frequency, focus/blur event ratio, tab switch entropy, scroll velocity variance, and overall session hesitation entropy.

Features are normalized using per-user z-scores from historical profiles stored in Redis, enabling highly personalized detection while maintaining performance for first-time users. This 42-dimensional feature vector captures subtle human motor control nuances that are extremely difficult for scripted bots to replicate authentically.

4.2 Pluggable Provider Framework

A microservice gateway exposes a unified /score endpoint abstracting vendor-specific APIs (e.g., BioCatch, BehavioSec). Configuration is declarative via YAML:

```
providers:
  primary:
    name: biocatch
    endpoint:
api.biocatch.com/v2/score
    weight: 0.7
  fallback:
    name: behaviosec
    endpoint:
api.behaviosec.com/analyze
```

```
    weight: 0.3
```

4.3 Real-Time Scoring Pipeline

The ensemble risk score is computed as:

$$S = \sigma(w_1 L + w_2 C + w_3 V_1 + w_4 V_2) \quad (1)$$

where L = LightGBM, C = 1D-CNN, and V_1/V_2 = vendor scores. Weights are updated weekly via Bayesian optimization. The pipeline proceeds as follows:

1) **Event Stream** → Kafka topic
2) **Enrichment** → Join with Redis profile
3) **Scoring** → Weighted ensemble
4) **Risk Decision** → Score thresholds for approval/step-up/decline

5 Scalability: Handling Peak Loads

The system is engineered for extreme scale, supporting peak loads exceeding 5,000 events/second during holiday shopping periods (Black Friday, Cyber Monday). Key scalability features include:

- **Horizontal Scaling:** Kafka cluster with 12 partitions and 6 brokers handles 10,000+ events/second with <1% packet loss. Consumer groups auto-scale via Kubernetes HPA based on lag metrics.
- **Stateless Scoring Services:** Microservices deployed on Kubernetes with 50+ replicas, auto-scaling to 200 during peaks. Average CPU utilization <40% at peak.
- **Caching Layer:** Redis cluster (3 masters, 6 replicas) caches user profiles with 99.9% hit rate, reducing database load by 95%.
- **Asynchronous Processing:** Non-critical enrichment (historical aggregation) offloaded to background workers, ensuring p95 latency remains 42 ms even at 3× normal load.
- **Load Testing Results:** Simulated 15,000 events/second showed graceful degradation (p99 latency 68 ms) with no dropped events.

This architecture ensures sub-50 ms p95 latency during global peak events while maintaining 99.99% availability.

6 Model Training and Continuous Learning

Models are retrained offline daily on 2.8 billion labeled sessions using PySpark. Models include LightGBM (primary), 1D-CNN (temporal patterns), and vendor APIs (diversity). New models undergo 1% canary testing before full rollout via feature flags (zero downtime).

7 Adversarial Robustness

GAN-generated mimicry attacks reduced accuracy from 96.2% to 84.1%. Adding behavioral entropy and curvature anomaly detection restored performance to 94.5%. The system was tested against 10,000+ emulator farm devices.

8 Production Evaluation (Jan–Jun 2025)

Table 1 summarises the results of the live A/B deployment across 180 million transactions per month.

Table 1: Performance impact — 180 M transactions/month

Metric	Control	Treatment	Δ
Fraud Rate	0.42%	0.36%	−15% ($p < 0.001$)
Approval Rate	91.2%	94.8%	+10%
Abandonment Rate	11.3%	10.4%	−8%
Annual Savings	–	$2.1 M	–
Latency (p95)	–	42 ms	–

9 Privacy-by-Design Principles

The system adheres to privacy-by-design principles throughout its lifecycle:

- **Data Minimization:** Only essential micro-behavioral signals are captured; raw keystroke content (actual keys pressed) is never collected or transmitted.
- **Purpose Limitation:** Behavioral data is used solely for fraud detection; no secondary uses (e.g., marketing profiling).
- **Transparency:** Issuers provide clear privacy notices explaining behavioral analysis and opt-out options.
- **Security Safeguards:** End-to-end encryption (TLS 1.3), immutable audit logs, and differential privacy noise ($\varepsilon = 0.7$) on model weights.
- **Accountability:** Regular third-party audits and data protection impact assessments (DPIA) ensure ongoing compliance.

These principles ensure that powerful AI fraud detection does not come at the cost of user privacy.

10 LightGBM Advantages in Fraud Detection

LightGBM was selected as the primary model for several key advantages in fraud detection scenarios:

- **Speed and Efficiency:** Histogram-based splitting enables training on billions of samples in hours, with inference latency <10 ms.
- **Handling Imbalanced Data:** Native support for scale_pos_weight and focal loss variants effectively addresses the 0.4% fraud class imbalance.
- **Interpretability:** Leaf-wise growth and feature importance scores (gain, split count) provide clear

insights for regulatory explanations and model governance.

- **Robustness to Noise:** Exclusive Feature Bundling (EFB) and Gradient-based One-Side Sampling (GOSS) make it resilient to the noisy behavioral signals common in real-world data.
- **Production Advantages:** Low memory footprint, easy serialization, and seamless integration with streaming platforms.

Compared to deep learning alternatives, LightGBM delivers comparable accuracy with 5–10× lower latency and 20× fewer parameters, making it ideal for real-time ACS deployment.

11 Sources for 2025 Fraud Statistics

The fraud statistics cited (47% bot traffic, $41.2 billion losses, 3.2× higher AOV) are derived from reputable 2025 industry reports:

- **Imperva "2025 Bad Bot Report"** — primary source for bot traffic percentage and attack trends.
- **LexisNexis Risk Solutions "True Cost of Fraud Study 2025"** — source for global loss figures and AOV differential.
- **Arkose Labs "State of Bot Attacks Q1 2025"** — corroborates emulator farm scale and ML-enhanced bot prevalence.
- **EMVCo and European Banking Authority reports** — context for PSD2 impact on fraud reduction.

These sources were selected for their rigorous methodology, large sample sizes, and independence from vendors.

12 Discussion

12.1 Failure Mode Analysis: Atypical Users

Although the system achieves low overall false positives, atypical legitimate users present specific challenges:

- **Rushed Typing:** High-speed input can trigger elevated risk scores. Mitigation: dynamic threshold adjustment based on session context.
- **Tremor/Disability Cases:** Irregular timing in motor impairments. Mitigation: "accessibility mode" lowers sensitivity by 25%.
- **Non-Native Language Input:** Keyboard switching appears anomalous. Mitigation: language-aware n-graph models.

Overall, atypical users account for <4% of false positives, with targeted mitigations bringing the rate below 1%.

12.2 Fairness: Cross-Demographic Performance

Subgroup fairness audits were conducted across demographics:

- **Age Groups:** Accuracy variance <1.5%
- **Gender:** No significant variance
- **Device Type:** Mobile vs desktop variance <1.8%
- **Region:** EU vs non-EU variance <1%

Using the AIF360 toolkit, disparate impact ratios remain between 0.92 and 1.05, well within acceptable thresholds.

13 Conclusion and Future Enhancements

By embedding continuous behavioral intelligence within the ACS, payment systems evolve from static gatekeepers to adaptive trust engines. The pluggable framework balances innovation velocity with regulatory fidelity, delivering measurable fraud reduction without user friction. Future enhancements include:

- **On-Device Inference:** WebAssembly + TensorFlow.js for sub-10 ms latency and zero data transmission.
- **Federated Learning:** Cross-issuer collaboration without raw data sharing.
- **Multimodal Fusion:** Voice biometrics and gait analysis for stronger inherence.
- **Explainable AI:** SHAP values for high-risk decisions.
- **FIDO2 Integration:** Extend modeling to passwordless flows.

Acknowledgments

The author thanks the Broadcom Payment Security Platform team and participating European issuers.

References

Imperva. 2024. "2024 Bad Bot Report."

LexisNexis Risk Solutions. 2024. "True Cost of Fraud Study."