Future of Finance:

How AI is advancing fraud detection in banking and financial services

Strategic technologies and implementation roadmap for financial institutions

COHERENT SOLUTIONS

Table of contents

Contributing experts:

Vittesh Sahni, Sr. Director of AI at Coherent Solutions

Shawn Torkelson, Chief Marketing & Strategy Officer at Coherent Solutions

Executive summary

Fraud scams are getting faster and smarter, while conventional rule-based systems are lagging. This is why banks are turning to Artificial Intelligence (AI). More than 70% of banks use AI to cut fraud losses by up to 50% and detect anomalies in milliseconds. Technology alone is insufficient. Trust is essential.

Customers expect seamless digital experiences, regulators require transparency, and businesses need to stay ahead of emerging threats. AI offers an all-in-one solution, provided it is used in conjunction with human-in-the-loop (HITL) monitoring, explainability, and robust governance.

This white paper examines how AI is revolutionizing fraud protection, including the technology underlying it, the real-world applications driving outcomes, and a roadmap for implementation that strikes a balance between innovation and regulatory compliance.

AI is becoming a strategic differentiator that leverages human judgement while preserving trust, compliance, and long-term customer retention.

Vittesh Sahni,
Sr. Director of AI at Coherent Solutions

Market context: Why fraud detection demands AI

The banking sector is undergoing a transformation as online channels grow and transaction volumes surge. As fraudsters utilize automation, synthetic identities, and deepfakes, the traditional "detect and react" model no longer works. What was once a manual investigation now requires intelligent systems that learn and act in real time.

AI-driven fraud detection continuously learns from transaction behavior, enabling institutions to process vast amounts of data instantly, identify complex fraud patterns, and react proactively. But this transformation is more than technological. True resilience in fraud management is a blend of AI precision combined with human judgement, a transparent audit trail, and strong governance.

As regulatory and competitive pressures grow, AI is evolving from an operational tool to a strategic risk management foundation that maintains financial integrity and strengthens consumer trust.

Rules vs. AI vs. hybrid systems

Feature / capability	Rule-based systems	AI-driven systems	Hybrid (Rules + AI)
Adaptability	Low – static rules must be manually updated	High – models learn from new pattern	Medium-High – AI adapts, rules cover known risks
False positives	High – rigid thresholds	Low – contextual, pattern-based	Moderate – balance between both
Speed & scalability	Limited – struggles with high-volume, real-time	High – supports millions of transactions per second	High – combines AI scale with rules as backup
Transparency	High – simple rules easy to explain	Medium – requires explainable AI frameworks	High – clear rules plus AI with monitoring
Implementation complexity	Low – easy to set up, but outdated quickly	High – requires data, training, and integration	Medium – incremental rollout possible
Effectiveness against new fraud	Low – predictable and bypassable	High – detects unknown, evolving fraud patterns	High adaptive with safety net rules

The takeaway is clear: *Rules are simple but ineffective, AI is powerful but requires governance, and hybrid systems combine the best of both worlds.* For most financial institutions, the hybrid model offers the right balance because it is adaptive enough to combat emerging fraud and transparent enough to satisfy regulators and customers.

Understanding AI for fraud detection

Fraud detection has undergone significant changes over the past decade. Modern defense relies on layered intelligence and human-in-the-loop (HITL) oversight to balance speed with transparency and accountability.

Core components

For many operators, consolidation is becoming the main path forward as mergers, acquisitions, and roll-ups are being used to reduce costs and compete more effectively. And, the next phase of value creation lies beyond U.S. borders for many operators and investors.

The European market, while large, remains fragmented. This opens opportunities for multinational operators to leverage proven models and build cross-border digital ecosystems that deliver consistent experiences while adapting to local markets and data privacy regulations. Investors also see potential for consolidation, and the organizations that unify experiences through digital platforms will gain a strategic advantage.

1. Supervised machine learning (ML) detects known patterns of fraud by learning from historical labeled data, making it effective when clear examples of fraudulent behavior are available. In contrast, Unsupervised models identify anomalies or unusual patterns without predefined rules or a target variable, making them especially valuable when labeled fraud examples are scarce or unavailable.

2. Graph analytics map fraud networks by revealing the hidden relationships between accounts, devices, and transactions.

3. Biometrics use physical and behavioral signals to prevent account takeover attempts.

System architecture

Data flows from payment networks, apps, and Know Your Customer (KYC) vendors into real-time pipelines. Features are extracted, a range of models' score events, and a decision engine responds in milliseconds. Human analysts review edge cases, creating a feedback loop that improves accuracy and reduces false positives.

The new frontier: Agentic AI

The second generation of agentic systems reason, plan, and act like analysts, linking transactions, modeling fraud scenarios, and generating explanations in natural language. These systems anticipate shifting fraud patterns and shift prevention from passive monitoring to active defense.

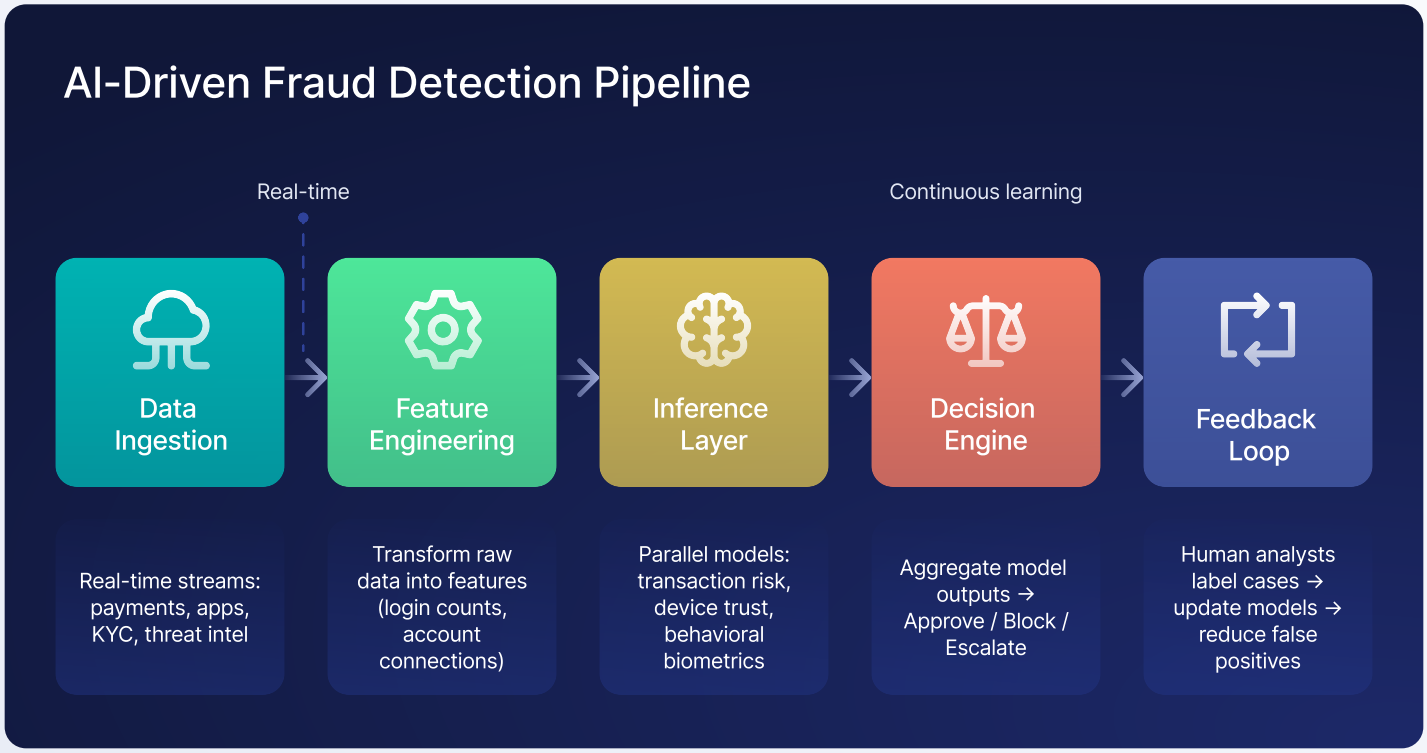

Agentic AI use cases & methods

COHERENT SOLUTIONS

AI revolutionizes fraud prevention across the customer life cycle, from onboarding to real-time payments. Its value materializes when ML velocity is paired with HITL oversight to ensure accuracy and accountability.

Some of the most powerful use cases in modern financial services are outlined below:

Real-time transaction monitoring

Live and historical transaction data are reviewed by AI models in milliseconds, and graph analytics uncover hidden fraud rings, while HITL review handles edge cases.

Success in the field: For a Midwest regional banking system, Coherent Solutions implemented a human-in-the-loop AI system to ensure compliance while minimizing friction during new user onboarding.

Dig deeper →

Document & KYC fraud

AI utilizes computer vision, deepfake detection, and pattern recognition to identify fake IDs and artificial selfies, flagging anomalies for human review.

Success in the field: Human-monitored AI ensures compliance without detracting from onboarding friction.

Dig deeper →

Biometric authentication

Liveness checks and facial recognition are increasingly used to combat AI-generated fake identities, while most platforms rely on built-in mobile biometrics (Face ID, etc.) for password less access.

Success in the field: Multi-modal AI authentication reduces account compromise by up to 70% and improves UX.

Dig deeper →

Crypto & digital assets

AI matches blockchain activity, clusters wallets, and applies natural language processing (NLP)to decipher dark web slang, signaling high-risk streams before losses escalate.

Success in the field: Coherent Solutions' AI-powered crypto monitoring analyzes blockchain activity, clusters wallets, and deciphers dark web slang using NLP to detect high-risk behavior early, preventing fraud and ensuring compliance to build platform trust.

Dig deeper →

Credit fraud

Machine learning, device fingerprinting, and NLP detect abnormal patterns and synthetic apps, reducing defaults and chargebacks.

Success in the field: AI reduces app fraud by 40–60% and lowers review costs.

Dig deeper →

Insider fraud

User and Entity Behavior Analytics (UEBA) tracks employee behavior, including increasing aberrant logins or suspicious data access, for human review.

Success in the field: AI+HITL enables early detection of insider fraud, ensuring trust.

Dig deeper →

Cross-border risk tracking

Graph AI and NLP enhance sanctions screening and high-risk corridor monitoring to address compliance needs with fewer false positives.

Success in the field: AI optimizes Anti-Money Laundering (AML) compliance efficiency by up to 35%, reducing authorized payment delays.

Dig deeper →

Agentic AI use case matrix

This matrix shows the direct link between financial risks, AI techniques, and business value. Each row pairs with a fraud scenario with the AI methods that address it and the resulting benefits. For example, graph analytics and machine learning enable the detection of anomalies in real-time, while AI-driven KYC and deepfake detection secure onboarding without adding friction. In short, the matrix illustrates how AI, combined with human oversight, enables a shift from reactive to proactive fraud management, thereby strengthening compliance, efficiency, and trust.

Use sase	AI techniques	Risk addressed	Digital value delivered
Real-time transaction monitoring	Anomaly detection, graph analytics, real-time scoring	Money movement fraud, account takeover	Fraud stopped instantly, smoother payments, higher trust
Document/KYC fraud	Computer vision, deepfake detection, pattern recognition	Synthetic identities, forged documents	Secure onboarding, regulatory compliance, fast user experience
Biometric authentication	Liveness detection, multi-modal biometrics	Spoofing, credential theft	Passwordless secure access, reduced friction, stronger security
Crypto fraud	Blockchain graph analysis, NLP monitoring	Laundering, scams, rug pulls	Safer crypto trading, FATF compliance, reduced losses
Credit fraud	Supervised ML, device fingerprinting, NLP	Application fraud, card misuse	Reduced defaults, faster approvals, operational efficiency
Insider fraud	UEBA, anomaly detection	Privilege abuse, data exfiltration	Early detection, protection of customer trust, accountability
Cross-border risk tracking	Graph AI, NLP sanctions checks	AML breaches, sanctions evasion	Faster, compliant global payments, reduced false positives

Customer experience and customer impact

Fraud prevention is no longer a back-office, after-the-fact affair; it now significantly impacts thecustomer's experience. Declined transactions, delayed onboarding, or unauthorized account holdsall erode consumer trust. However, AI-driven fraud detection can simultaneously enhance security and consumer satisfaction.

AI enhances the speed and accuracy of fraud detection, minimizing false positives, and enabling legitimate transactions to proceed smoothly. Real-time monitoring and biometric authentication offer frictionless access, building trust, and loyalty. Human authorization validates high-risk decisions, striking a balance between efficiency and accountability.

AI also enables personalized, context-aware risk management, adapting protections based on individual user behavior. Customer benefits are:

- Reduced false declines and disruptions, which leads to less frustration.
- Greater trust resulting from proactive security.
- Increased loyalty and engagement as frictionless, no-risk experiences foster long-term use.

Well-implemented, AI turns preventing fraud into a strategic differentiator, not a friction point, driving satisfaction, loyalty, and lifetime value.

Strategic business benefits

Incorporating AI into fraud detection provides financial institutions with a competitive edge by leveraging machine learning, coupled with human oversight, to harmonize automation with judgment and expertise.

The results include:

1. Improved accuracy: Analyze vast volumes of data to detect sophisticated patterns of fraud, reducing false positives, while enabling legitimate transactions to clear without hindrance, supporting customer trust. For instance, AI-based AML solutions can detect 2–4 times more suspicious activity while reducing alert volumes by over 60%.

2. Real-time detection: Continuously monitor user activity and transactions in real-time, enabling an immediate response. Applications like BioCatch analyze tens of thousands of behavioral and device indicators in real-time, minimizing losses and enhancing the customer experience.

3. Enhanced user experience: Eliminate unwanted interruptions and false positives, allowing seamless digital experiences while keeping fraudsters at bay.

4. Lowered operating expenses: Handle repetitive detection, forwarding only edge cases to human analysts. This HITL approach lowers investigation costs without compromising security.

5. Compliance with regulations: Meet regulatory demands for transparency and accountability, thereby reducing legal risk and strengthening customer and regulatorconfidence.

6

Business agility: Adapt to unique fraud patterns through feedback mechanisms and human intervention, protecting revenue and maintaining a competitive edge.

7

Reduced direct financial losses: Proactively prevent fraudulent transactions, account takeovers, and chargebacks, significantly cutting direct monetary losses, reimbursement expenses, and operational investigation costs. Early detection minimizes the financial burden of fraud and helps preserve revenue integrity.

8

Improved customer retention: By reducing false positives and eliminating unnecessary friction, customers experience smoother, more secure interactions. Positive digital experiences build trust, reduce churn, and encourage long-term engagement across products and services.

9

Strengthened company reputation: Demonstrate a high standard of security, transparency, and reliability. Effective fraud prevention not only reduces risk but also reinforces customer confidence, strengthens relationships with partners and regulators, and enhances overall brand credibility.

With the aid of AI-driven fraud prevention, institutions can enhance security, operate more effectively, improve customer experience, remain compliant, and achieve strategic business advantages.

Compliance, risk & technical challenges

AI enhances fraud detection, but it also brings regulatory, ethical, and technical challenges. Financial institutions must strike a balance between innovation, compliance, and trust while addressing risks.

Privacy & data protection:

Fraud models rely on personal and behavioral data, which relate to the GDPR, AI Act, and other privacy regulations. Key risks involve obtaining explicit consent for biometric and behavioral data, balancing data minimization with the need for high model accuracy, managing cross-border data transfers, and addressing the technical challenges of removing personal data from trained models to uphold the right to be forgotten.
The EU AI Act outlines AI classifications for biometric identification and creditworthiness, such as high-risk, mandating risk assessments, continuous monitoring, robust logging, and human oversight.

As above, you go right into the subsections and areas of concern but and maybe it will be ok in the whitepaper but it might make sense to add a lead-in sentence. Just specifying... some of the key challenges include... — as you can see i like the word 'include' :)

Explainability & regulatory expectations:

Complex AI models, such as deep learning and graph networks, achieve high accuracy but often lack interpretability. This makes it difficult to explain automated decisions to regulators and customers, and autonomous rejection of transactions without human review can breach accountability and transparency requirements.

Fairness & bias:

AI systems may unintentionally amplify systemic bias due to historical imbalances in training data, proxy variables like ZIP codes that do not adequately reflect demographic patterns, or feedback loops that reinforce false positives. Regulators increasingly require bias assessments, mitigation strategies, and transparency (see: FTC, EU AI Act, FDIC/OCC guidance).

Integration & technical challenges:

Deploying AI within financial institutions is complicated by legacy infrastructure and fragmented data environments. Effective real-time fraud detection relies on low-latency access to clean, consistent data, modern feature stores and model-serving frameworks, and adherence to governance, versioning, and audit requirements.

Monitoring & model maintenance:

AI models degrade over time because of concept drift and pipeline errors. Institutions must maintain continuous performance monitoring, schedule regular retraining and validation, and preserve audit trails and version control to comply with regulatory expectations from bodies such as the SEC and FDIC.

Emerging AI use cases:

New agentic and generative AI applications, such as autonomous investigations, real-time identity verification, and synthetic data generation, introduce novel risks. These include autonomous decision-making without sufficient human oversight, vulnerabilities in access control and adversarial resilience, and the ongoing need for independent validation and strong governance mechanisms.

FINRA & reporting obligations

AI-driven fraud detection must comply with FINRA oversight and reporting standards to ensure transparency and accountability. Key points to ensure compliance include reporting fraud or technology disruptions with market impact, maintaining strong documentation, validation, testing, and version control for AI systems, and aligning surveillance tools with market integrity and customer protection standards. Firms should also enable auditability and supervisory review for all AI-driven decisions and oversee third-party AI tools to ensure compliance with FINRA requirements on outsourcing, cybersecurity, and data privacy. Compliance with FINRA ensures accountable, auditable, and transparent AI fraud defense operations.

The future of fraud defense

Financial fraud is becoming increasingly sophisticated, outwitting traditional rule-based and static AI approaches. The future of fraud prevention lies in adaptive, human-in-the-loop AI systems that continuously learn from real-world transactions, behaviors, and network activity to detect new fraud techniques in real time.

While big data once gave companies a clear edge, today the real challenge lies in navigating privacy restrictions and the shortage of high-quality, labeled fraud data.

Synthetic data generation and federated learning are methods that enable scalable training without exposing sensitive or proprietary live data, ensures high-quality model training and testing.

Independent AI programs can now spot, score, and even respond to suspicious behavior, yet high-impact actions, such as freezing accounts or reporting fraud rings, continue to require human judgment. Adding explainability to such programs offers transparency and accountability to regulators, auditors, and consumers.

AI-driven fraud prevention with human oversight delivers measurable benefits. To achieve these benefits, institutions need to:

1. Deploy adaptive, HITL AI to balance automation with accountability.
2. Invest in privacy-preserving data strategies, such as federated learning and synthetic data.
3. Create explainable models for transparency and trust.
4. Retrain and continuously monitor models to stay ahead of new fraud patterns.
5. Use rapid prototyping to advance solutions before fully deploying them.

Coherent Solutions, Inc.
1600 Utica Ave. S., Suite 120
Minneapolis, MN 55416.
+1 (844) 224-4994

Clutch
LinkedIn
www.coherentsolutions.com

Contributing experts:

Vittesh Sahni, Sr. Director of AI at Coherent Solutions
Linkedin

Shawn Torkelson, Chief Marketing & Strategy Officer at Coherent Solutions
Linkedin

COHERENT
SOLUTIONS