IJHSR
ijhighschoolresearch.org
REVIEW ARTICLE

# AI for Fraud Detection: A Cross-Industry Analysis of Banking and E-Commerce

Andy Wu

Cypress Bay High School, 19113 S Hibiscus St, Weston, Florida, 33332, USA; awesomeandy351@gmail.com
Mentor: Dr. Siddharth Krishnan, Samuel Lefcourt

ABSTRACT: Banking scams and e-commerce fraud steal billions of dollars from consumers and companies each year, fueled by identity theft, payment fraud, and deceptive methods that exploit weak verification systems. E-commerce fraud, on the other hand, takes the form of misleading reviews, fake storefronts, and unauthorized or fraudulent transactions. They tend to blend more easily into the decentralized world of online businesses, making it harder to spot and verify fraudulent activity. Artificial intelligence offers a simple solution. By analyzing patterns in customer behavior and transaction data, AI can detect fraud in real time. This review considers whether AI strategies employed in the banking sector will be similarly functional in online shopping and the general flexibility of the AI strategies in varied environments. While results demonstrate that AI performs well across industries, adjustments are necessary due to varying regulatory frameworks, data availability, and consumer privacy standards. This paper addresses gaps in the literature regarding cross-sector AI strategies, proposing that hybrid models and shared data infrastructures could enhance risk management and trust. As AI tools advance, this research suggests that scalable, transparent systems for fraud detection will become required for global financial fairness and integrity.

KEYWORDS: Systems Software, Cybersecurity, Artificial Intelligence, Banking & E-commerce, Fraud Detection.

## Introduction

Financial fraud is a multi-billion-dollar problem. While AI has become an essential tool in real-time fraud detection, simply moving an AI model from one industry to another is fraught with complexities. This paper explores the unique challenges of applying banking-based AI strategies to e-commerce.

A scammy snapshot of reported U.S. fraud losses.

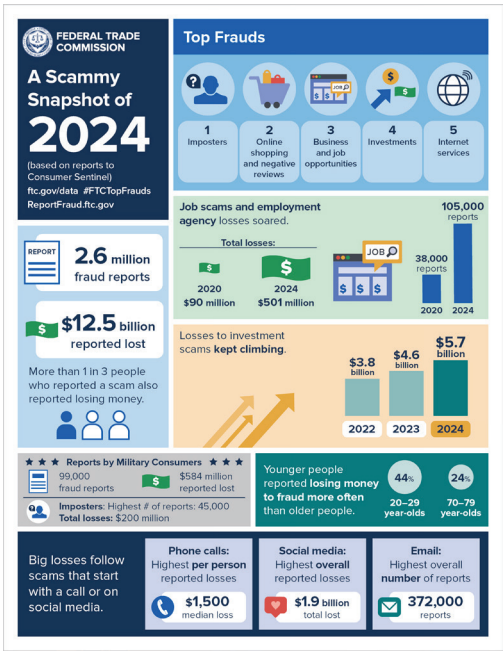


**Figure 1:** This infographic, produced by the Federal Trade Commission, illustrates major fraud trends that were reported in 2024 and nearly all the main scams, impostors, online shopping fraud, and investment scams, and how much money they cost, adding up to $12.5 billion in losses.

Figure 1 shows the impactful increases in job scams and investment scams, how military consumers continue to be impacted, and how fraud is impacted by factors such as the method of contact and age. Overall, the infographic tells us how scams are changing and who is hit the hardest. Most scams, such as impersonation, online shopping fraud, false business opportunities, and other scams, affected people of all ages and occupations. Young adults, in particular, were more likely than older adults to lose money to fraud,[1] highlighting just how prevalent and personalized these threats have developed in society today! Empirical research demonstrates that individuals aged 18-34 report online fraud victimization at rates 2.5 times higher than those over 55, attributed to greater platform engagement, adoption of emerging payment technologies, and higher risk tolerance in digital transactions.[2] Behavioral economics studies reveal that younger consumers exhibit present bias and optimism bias that reduce fraud vigilance, while older adults' lower digital literacy paradoxically correlates with more conservative online behavior that limits exposure.[3] The good news is that advancements in artificial intelligence (AI) are also delivering solutions for us to fight back. By using customer behaviors and payment patterns, AI can actually intercept fraudulent payments as they are happening.

This paper raises an important question: can we apply the AI tools developed for the banking and financial services sectors to online shopping? This question is important because it is valuable to know if banking-based AI models are transferable to e-commerce. After all, fraud schemes commonly cross sectors, and the use of shared solutions can limit duplication and enhance efficiency. As digital transaction models are beginning to unify so that you cannot distinguish between financial ser-



DOI: 10.36838/IJHSR89.237

vices and retail platforms, unified solutions could have a larger net through shared risk from fraud events. If there is a strong belief that established banking-based models can be leveraged to develop viable fraud detection solutions for online shopping, the adoption of the unique context of the retail domain could also expedite the development of scalable fraud detection systems, particularly where applicable verification systems are less strong. This paper also examines the extent to which these systems can adapt to different environments. There has been a substantial amount of research conducted examining all of the different ways machine learning, deep learning, and transfer learning can be useful in fraud detection. Several studies have included applications related to the tracking of credit card abuse,[4,5] the detection of spam and fake reviews,[6] and analyzing risk and fraud in online payments.[7] These studies concluded that AI is effective in protecting against fraud and related activities and that it can be applied to both banking and e-commerce.

However, the recognition that adjustments and modifications had to be made to algorithms and technology was also important. The focused and relevant challenges include differences in legislation and regulations, privacy and consumer protection laws, and regulations for the data available in one industry not being available in another industry. For instance, banks are often subject to strict regulations, have consistency in the data set, and have mostly reliable data.[8,9] Conversely, online shopping marketplaces have less reliable data and more privacy issues.[10]

This review articulates impactful gaps in existing research. Very few pieces of research mention sectoral sharing of fraud detection systems. While it is easy to believe a model trained on banking data could be simply reused in e-commerce, this is fraught with complexity. This paper argues that hybrid AI models and sharing data platforms can play a significant role in detecting fraud, building trust, and decreasing fraud loss in both banking and e-commerce.[11, 12] This, in paired with the advancing power and widespread use of AI, justifies the development of nationally and globally scalable fraud systems that are clear and fair for everyone involved, especially to maintain financial systems worldwide.

## ■ Discussion

This article explores the pros and cons of AI-powered fraud detection in banks and e-commerce businesses. In general, AI technology is very efficient in a well-organized and regulated banking environment, but its application in e-commerce is limited to a certain extent. This is because user behavior varies greatly, datasets are fragmented, and there are hardly any verification steps in e-commerce. As a result of differences in regulatory frameworks, data availability, and transaction patterns, models cannot merely be transferred from one sector to another without proper retraining and feature adjustment.

Combining model structures and transfer learning has the potential to solve this problem by letting the banking-focused AI models' components guide e-commerce detection with the help of domain-specific adaptations. Federated learning may also facilitate cross-platform collaborations by solving the problem of data fragmentation while still complying with privacy standards. However, several obstacles, such as algorithmic bias, privacy laws, and data siloing that are caused by practical, legal, and ethical issues, still exist.

The use of AI across different sectors for fraud detection is one of the main ideas left over from the discussion. However, for this to work, a context-aware model, real-time behavioral analytics, and collaborative intelligence-sharing frameworks are necessary. Research made in the future has to be on the performance of longitudinal models, scalable hybrid systems, and regulatory and organizational factors, so as to ensure that AI systems are still efficient, fair, and flexible in different financial ecosystems.

### *Common Fraud Techniques and AI Basics:*

Fraud often refers to an individual or personal entity manipulating a system to steal funds and/or personal data. In banking, fraud could mean, among other things, stealing someone's credit card information, opening a fraudulent bank account, or phishing someone into accessing their funds.

Cybercriminals can extract money from unsuspecting individuals through card skimmers, illegitimate credentials, and social engineering.[13]

Artificial intelligence is now being used to combat some of these tactics. In terms of banking, machine learning models can detect user behavior and money transfers that don't coincide with how the typical person spends their funds, making it easier to catch fraud.[6] Deep learning technologies, specifically transformer-based models, provide the capability to detect rare events for fraud detection, unlike older technologies.[10] These models are trained using large datasets of past transactions and login behaviors that enable them to identify what looks "normal" and raise flags for items that fall outside of the norm. For example, a transformer model can outperform other models by taking into account the order of events, such as a new log-in from another country immediately followed by a high-dollar wire transfer to detect complex multi-step fraud, which other models have not been capable of effectively detecting.

AI is used in e-commerce to help track and detect consumer actions, track bot traffic, or to authenticate that a user is a human being for user verification using biometric recognition systems that help curb fraudulent accounts or card testing attacks.[7] Other systems use federated learning and teach companies how to exchange information about fraudulent behaviors without having to provide all the information, which is another advantage of secure flaws and fraud detection improvement.[14] If we formalized federated learning, we would describe it as a decentralized machine learning mechanism whereby models are trained locally, on user devices or servers, and only the model updates are shared without raw data. This enables shared learning while minimizing the impact of sharing sensitive information, which aids in privacy.

Within the past few years, the landscape of using AI for fraud detection has evolved from rule-based systems to machine learning and deep learning models capable of detecting advanced, evolving fraud patterns. While there are many different variations of algorithms in AI systems used for fraud

detection, each set of algorithms has its own specialty in detecting financial anomalies. Traditional machine learning algorithms are still the most basic form of fraud detection. Logistic Regression is the most widely used machine learning algorithm and is suitable for most binary classification tasks, such as whether a transaction is fraudulent or not, based on features including amount, time, and location.[15] Decision Trees and Random Forests are more sophisticated classifiers in that they model interactions between features and aggregate multiple decision trees.[16] The decision tree style of classification is very effective for structured financial data sets and has been shown to be superior to naive classifiers in banking and related applications. Support Vector Machines are generally the go-to model for anomaly detection, since they define the optimal border between what is normal compared to suspicious (when working with a more feature-rich data set).[17]

As fraud schemes continue to evolve, deep learning models can be the most effective solutions. Recurrent Neural Networks (RNN) and Long Short-Term Memory (LSTM) networks are very effective for sequential data, such as login histories or transaction timelines, which capture inherent temporal dependencies and detect multi-step fraud schemes that traditional modeling approaches would miss.[18] For example, LSTMs were able to identify sequences of behavior in which a user appeared in a different geographical location (e.g., their residence) followed by high-value transfers. Autoencoders, which are a type of unsupervised learning, are also a powerful way to identify anomalies. They learn the structure of normal transactions and flag anomalies to identify instances of prior unencountered forms of fraud that do not require training on labeled data.[19]

Fraud can arise from multiple interconnected entities, forming networks of financial crimes. Graph Neural Networks (GNN) have inherent relationships within them and are specifically designed to study these relationships. Understanding behavior in user-device-transaction graphs can also help detect fraud rings.[20] GNNs have been shown to successfully identify collusive behaviors and coordinated attacks among behaviors occurring in a financial network. Natural Language Processing (NLP) models like BERT(Bidirectional Encoder Representations from Transformers) can be helpful to detect fake reviews and fraudulent communications. LSTM models can be fine-tuned and applied to extract linguistic features that identify an unusual number of inconsistencies in sentiment behaviors from user-generated content.

These technologies reveal that even though fraud is evident in the banking environment, while another in the online shopping environment, AI can adjust to both environments to help protect people and organizations. However, the context in which AI operates, how data is collected, what rules govern its use, and how predictable user behavior is, play a major role in determining how well it performs. This sets the stage for the next sections, which explore how AI works in banking, how it struggles in e-commerce, and whether fraud detection systems can be reused across these sectors.

***AI in Detecting Fraud in E-Commerce:***

The e-commerce environment is inherently less controlled than banking, and this chaos amplifies the challenges for AI-driven fraud detection. Unlike a bank, where accounts are linked to verified identities and transactions occur in a tightly regulated framework, marketplaces such as Amazon, eBay, and Alibaba operate as sprawling ecosystems of independent sellers. Many have minimal verification requirements for opening an account; some require nothing more than an email address and a linked payment method. The low barrier to entry creates opportunities not only for genuine entrepreneurs but also for bad actors who can take advantage of the system's openness; data trustworthiness is the most significant obstacle. While transaction records for banks are sequential and therefore simple to format into structured data for AI models, e-commerce is much messier; product records can be begrudgingly complete, inconsistently labeled, or sometimes completely fabricated. Merchants may take a product image and description from someone else's listing, manipulate a price history to give the impression of a legitimate discount, or maintain multiple accounts to avoid account bans. Customers also may go unverified, making it nearly impossible to track bad buying patterns across merchants and platforms. Unlike banking systems, where regulatory compliance mandates identity verification, e-commerce platforms prioritize frictionless checkout experiences that enable anonymous or pseudonymous transactions.[21] Research analyzing over 2.3 million transactions across multiple platforms found that mandatory authentication reduced checkout completion rates by 18-34% but decreased fraud attempts by 41%, illustrating the fundamental tension between user experience optimization and security in online retail environments.[22] This structural difference creates asymmetric information problems where platforms lack the verified identity foundations that enable effective fraud pattern recognition across customer lifecycles, a challenge absent in regulated banking contexts.

AI detects and prevents fraud in e-commerce, specifically, through analyzing user behavior data, transaction data, and biometric data.[23] This infographic shows how a machine learning model can raise a red flag for anomalies in transaction data (fake reviews, bot purchases, or false purchases) using many risk factors of comprehension of user behavior. Further to this, it illustrates the frustration of fragmented data across different platforms; ultimately, AI will be only as good as the quality and completeness of the data that drives the decision trees. This infographic illustrates just how far AI has come in managing that type of complexity.

These machine learning models often can detect patterns of fake reviews based on observable intensity and frequency discrepancies in the language of the reviews. So if we see 5-star ratings for a relatively unknown seller come in quick succession, those should raise a red flag for further attention. Bot detection systems identify automated scripts by analyzing engagement behaviors. These systems distinguish bots from human users by examining interaction patterns: bots exhibit consistent timing and linear navigation paths, while humans show variable speeds and organic browsing behavior.


DOI: 10.36838/IJHSR89.237

The model will conduct an analysis of the known patterns of browsing behavior for a human, and compare those engagements with query results and speed returned by a script.

Purchase-risk scoring systems can evaluate dozens of signals at once - for example, from anomalies in shipping address, to what might reasonably be considered unique combinations of ordered items - to support operatives in determining if transactions are unlikely to be fraudulent, before the completion of the transaction.[11,13] AI might consider a mixture of the risks of product, the order size, and the history of addresses to flag an order for five high-end smartphones, shipping to a residential address that has not had a parcel delivered therein before this one.

Some platforms do even further by combining biometric identity verification measurement (for example, facial recognition when a seller account is created), or behavioral logging of a user's regular path of navigation, to measure normality versus deviation of behavior. But unfortunately, all of this remains in independent companies. Privacy acts such as the GDPR and competition laws applicable to, or establishing jurisdiction, prevent platforms from explicitly sharing data; therefore, AI models only have a one-platform view of user behavior. In addition, no one platform helps ensure that a seller, who was banned from a platform, does not just restart their activities with another platform, again creating the same situation without any inter-platform knowledge/ intelligence of this actor's previous work. This is a major disconnect and is one of the most common weaknesses of fraud prevention measures that are being applied to e-commerce. AI, for all its pattern recognition power, is only as strong as the completeness of its input, and right now, those inputs are fragmented.

Fraudsters take advantage of these systemic vulnerabilities by means of what researchers call "platform arbitrage"; they are basically simultaneously testing their ability to avoid detection on various platforms before they can launch their coordinated attacks at full scale.[24] An empirical study of the cross-platform fraud networks discloses that the sellers who have been banned reappear on alternative marketplaces within median timeframes of 12 to 15 days, and in most cases, they initiate their activities by utilizing the same business models and the same supplier relationships that have not been recognized due to information silos.[25] In such a scenario, AI models are only able to see through a limited perspective. The same actor can be engaged in the perpetration of a particular fraud on several platforms without the possibility that the respective platforms will receive an alert, since there is no insight into the patterns that go beyond the platforms.

***Determining the Limitations of Artificial Intelligence in E-Commerce:***

The existence of isolated or siloed e-commerce systems exacerbates these challenges. Platforms work independently, keep their proprietary datasets in closed silos, and seldom try tackling fraud together. That is understandable from a business perspective, as no company wants to disclose sensitive transaction data or badge their unique selling proposition to competitors, but it has a cost. Fraudsters delight in these failures, as they pair knowledge of one platform's whitespace with yet another one. In such a setting, AI models only see through a narrow lens. A given actor may perpetrate the same fraud across multiple platforms without any alerts, as there is no visibility into patterns that traverse both platforms. For example, an anonymous buyer could complete multiple chargebacks on one site and simply switch the whole process to another marketplace and an alias. This happens with no shared threat intelligence for everyone, and so it does not constitute collaboration, and it looks episodic rather than coordinated.[12,14]

Another obstacle is the vast amount of noise in the information that informs these AI models. For example, sellers can copy-paste similar product descriptions from competitors, which can skew their similarity metrics to non-fraudulent buyers. Fraudulent buyers can work their way through different IP addresses or payment methods to mask their identity. Honest users can also behave in random, unexpected ways, as might occur if the buyer procures a bulk order of an unusual mix of products, which can create false positives that diminish trust in the systems we want to build. Behavioral studies document that life transitions (residential moves, marriages, births) generate purchasing patterns statistically indistinguishable from fraud indicators: analysis of 1.8 million transactions found that consumers experiencing major life events exhibit 67% higher transaction variance, 43% more frequent address changes, and 2.3 times greater product category diversity compared to stable-state purchasing.[26] Gift-giving behavior similarly produces transaction profiles overlapping with fraud signatures, multiple shipping addresses, cross-category purchases, and elevated transaction values, requiring context-aware temporal models to distinguish seasonal legitimate activity from coordinated attacks.[27] This unpredictability stands in stark contrast to banking environments, where the activity associated with an account is usually relatively stable over time.

Behavior tracking and biometrics can also provide better datasets, but their data is limited when confined to a single platform. For example, a seller can verify their identity with minimal subsequent checks, and can then relinquish control of their operations to an associate or employee who intends to commit fraud, and those safeguards can be undone. And while biometrics at least accurately define the accountholder as an individual, they don't eliminate the likelihood that the monitored individual would engage in purposeful fraud.

As the research found, while it is possible to deploy an AI system trained in one marketplace environment, such as cancellation rates or fraudulent attempts reporting, in another environment without retraining the model, you are likely to lose at least some precision in your reporting.[30] For many different reasons, including changing demographics, purchasing behaviors, and product categories, you run the risk of deploying a model that has emerged from a high-volume electronics marketplace and mistakenly classifying completely normal activity as fraudulent in a niche craft marketplace and vice versa.

In short, detached e-commerce AI is good at observing and capturing a local event, but has serious limitations with regard to identifying and connecting global events. While algorithmic improvements remain relevant, the central argument empha-

sizes the need for new frameworks that enable collaborative fraud intelligence across platforms, without compromising privacy regulations or competitive boundaries.

***Methods Used Between Domains in a Commercial Setting:***

The issue of whether banking fraud detection models can be repurposed to e-commerce is more than just a technical compatibility issue. Banking data is very structured; every transaction has consistent summary details (amount, timestamp, merchant code, and account metadata). E-commerce and, therefore, e-commerce fraud detection data are ultimately transactional, but they are more complex: a single transaction may include many items, sellers may have their own ID systems, and many/multiple products could be shipped internationally with non-standardized address formats.

Transfer learning[9] can be posited as a theoretical link between these "environments". In essence, it allows an AI model to take the knowledge it learned from one dataset, say, detecting suspicious patterns in financial transactions, and adapt that knowledge to another task, such as flagging abnormal purchasing behavior online. The primary benefit is efficiency: the model could potentially make use of prior information about the domain, saving training time and computational costs. However, some practical obstacles come to the forefront.

Generally speaking, fraud in banking often manifests itself as stark changes from an individual's prior behavior, for instance, a $3,000 transaction in a foreign country. However, high-variance behavior is inevitably present in e-commerce, yet not in fraudulent forms, for example, holidays, flurry buying as a result of flash sales. Time-series analysis of consumer purchasing patterns reveals coefficient of variation (CV) values 3.8 times higher in e-commerce contexts compared to banking transactions, with seasonal promotions generating transaction volume spikes of 300-450% and individual-level behavioral volatility that confounds standard anomaly detection thresholds.[28] Furthermore, legitimate cross-border shopping creates geographic transaction patterns that banking fraud models flag as high-risk: studies show 23% of online shoppers regularly purchase from international merchants, generating location-based anomalies that represent normal global commerce rather than account compromise.[29] This differs from the case presented by banking, where a model would potentially need to be retrained and/or at least fine-tuned to detect the different basic behavioral and risk signals.

The second practical barrier lies in the way and the features each model weighs. An e-commerce model would need to weigh sentiment analysis, order frequency, and geographic clusters for shipping, whereas a banking model may weight merchant codes or poorly studied ATM withdrawal patterns significantly. If I am simply taking one model into another space without adjusting the features pertinent to that model, I am only inviting misclassification issues.

Furthermore, there are also legal and technical barriers that bar the portability of transaction data. Privacy laws generally prohibit the raw transaction data from leaving the sector it originated from, meaning any export must occur under either anonymized, pre-processed, or synthetic datasets. For some use cases, these may be more than acceptable, but not necessarily suitable to buffer from retraining, which could be impeded by the lack of independence and potentially dilute the value of transfer learning. Researchers[1,6] propose a hybrid of transfer learning, using transfer learning where the architecture generalizes well (for example, in an anomaly detection architecture) to develop sector-specific components for the environment's challenges. This way, researchers can get the benefit of both approaches (e.g., avoiding starting over inefficiently while still developing to the unique requirements of specific domains).

Hybrid architectures are a viable option to keep components that can be transferred, like anomaly detection frameworks, while at the same time, they incorporate domain-specific layers that are tailored to each environment.[9] For example, banking implementations focus on deviation-based alerts and transaction history analysis, whereas e-commerce versions use seller reputation scoring, review sentiment analysis, and seasonal pattern recognition that allows purchasing variance.[10,15] This modular method lessens the redundant development and keeps transfer learning efficient while retaining context-specific precision. Research suggests hybrid models can achieve almost the same level of accuracy as fully customized systems, and at the same time, they require a lot less training data and development time, which makes them viable for cross-sector organizations or platforms with limited resources.[9] The main problem is figuring out the best points at which the components should be shared and the ones that should be specialized.

The most effective AI fraud detection systems in the real world are highly customized to the environments in which they operate. Amazon's detection system continuously consumes and analyzes enormous amounts of behavioral data - how quickly a user transitions from product page to checkout, whether shipping addresses fit the patterns of known fraud, and how review content compares to normal lenses of human language. The AI models can surface anomalous behavior in milliseconds, which can hold orders for human review before the confirmation of payment.[31] For example, their models are trained to detect "friendly fraud," which occurs when users have fundamentally higher rates of "item not received" claims in comparison to their geographic peers (or with other purchase histories).

EBay approaches analytics differently. They incorporate both buyer and seller perspectives, using a blend of data points to create dynamic trust scores based on transaction histories, behavior anomalies, and dispute rates. For example, high trust ratings may set up manual reviews of many high-value listings by sellers with a spiking amount of activity. eBay may also take a closer look at buyers with new accounts who are making unusual bulk purchases.

Mastercard's fraud protective systems are on a global scale, reviewing billions of card transactions every day, currently and in the future. Their model incorporates both supervised learning, based on previous fraud alerts, and unsupervised anomaly detection before authorizing payment. Their continuous updating and direct deployments are more like an in-real-time approach to daily complexity, enabling them to rapidly adapt to a new pattern of fraud.[30] Reports suggest that these systems


DOI: 10.36838/IJHSR89.237

can achieve precision rates of above 99% and a false positive rate that is less than a fraction of a percent of transactions.[32]

What these systems all have in common is a feedback loop: confirmed fraud cases go back into the model, allowing it to learn from its mistakes as well as its successes. Having a proprietary, high-quality dataset, usually compiled over years, is arguably the biggest advantage, as these companies consistently maintain their precision rates. However, this advantage is also the reason that portability is limited. For example, Amazon's system is designed for e-commerce shopping carts and shopping reviews. Mastercard's models' outputs do well with structured banking transactions, but not directly in shopping transactions.

Any attempt to adapt one of the other systems for the other would require some fundamental architectural changes as well as retraining for a considerable period of time. While the technology can travel, the models do not.

***Limitations, Ethics, and Future Directions:***

Even with progress, AI fraud detection can have blind spots. New fraud techniques produce limited amounts of data; therefore, even the most sophisticated systems will lag behind agile adversaries. In practice, fraud rings use "test runs" - low-value transactional tests to probe for weaknesses in the system. If the "test runs" are successful, they can be deployed at scale and with coordination, having seen the system fail. User behavior is another important consideration. High-volume transactional events, such as Black Friday or periods of commercial online transaction spikes from pandemics, may develop behavioral signals that may trigger alerts in velocity or volume, but are normal transactions. AI systems that do not take into consideration time sensitivity and adaptive thresholds will quickly overwhelm the review effort by human reviewers if they even trigger, and then bottleneck the transaction by interjecting a fear factor. Ultimately, this will threaten the validity of AI predictions if human intervention becomes overwhelming and revisionist.

Additionally, an important performance constraint that is discipline-dependent but often ignored is algorithmic bias. Models are trained on a historic dataset that can perpetuate historic bias. For example, if the historic fraud profile is overrepresented in a specific geographic location, the AI will over-flag all transactions received from that specific geographic location and create a discriminatory feedback loop. This has serious ethical implications and will exclude segments of consumers. Aside from the ethical implications, auditing models for fairness and using balanced and representative training datasets are just as good ethical behaviors as they are data obligations.

Some types of fraud are also difficult to identify with transaction data alone. Social engineering fraud is described as criminals convincing the victims to do the work of initiating the transfer or disclosing credentials. These criminals are acting to circumvent the automated fraud detection measures that only assess behavior changes from the record of the user's typical behavior pattern. Deepfake-based identity theft goes even further, as synthetic audio and video generated from deepfake systems can undermine the trust of biometric verification systems by accurately depicting the legitimate user's face and voice.[7,23] In order to sufficiently mitigate these threats, organizations will need to use AI models that employ multimodal verification approaches to verify identity by distinct signals in real-time.

At the systems level, the fragmented nature of data across organizations will continue to be a hurdle to achieving coherent fraud control. Unless there are trusted, privacy-oriented solutions that allow for intrinsic sharing intelligence, organizations will have to continue to prepare their defensive strategies on an isolated basis. Federated learning is a viable approach to circumvent the data fragmentation challenge. Each organization's machine learning model is trained locally on their own data - only the learned characteristics and features, which capture the mental representation of the data for the training model, are shared. This enables collaborative improvement without exposing sensitive information.[5] However, a federated model may encounter technical/ governance issues when it comes to ensuring a review process of integrity updates, or even the risk of embedding bias from one partner regarding poor data, and ensuring multi-jurisdictional compliance.

Looking ahead, three trends are more obvious: Hybrid models that involve supervised learning, anomaly detection, and knowledge graphs that map the relationships of connections between all the entities, such as users, devices, IP addresses, and concealed systems of fraud, may improve resiliency to various forms of unknown fraud. No one should build a template fraud detection model. Predictable phenomenological and congruous forms are outside of predictability, and therefore, the collection of time-based dynamic scores via behavioral biometrics, keystroke dynamics, motion-clicking actions of engagement, or touch screen affect can facilitate verification without authorization. Explainable AI (XAI) structures may help with regulators while providing user acceptance of a hybrid, with expressly clear, useful, verifiable clarifications about why a shown transaction could be flagged. International

regulatory charities help to enable cross-platform intelligence-sharing, noting they must respect the law in terms of data privacy. How developments in this area play out one day, the goal of transferable fraud detection across banking and e-commerce may go from a dreamed possibility of the enthusiast to an industrialized baseline standard.

## ■ Conclusion

Fraud detection in the banking sector and in e-commerce may have similar purposes, stopping criminals before they can commit harm, but the ecosystems of each sector are different. Banking is structured with strict regulations, fixed standard data formats, and customers with fixed identities.

E-commerce works in a high-variance, loosely monitored environment with anonymity, and where data quality is never sequential nor uniform. Differences such as these illustrate why an AI model that achieves success in the banking sector cannot merely be dropped into an e-commerce environment.

Transfer learning will provide a helpful bridge between the financial and e-commerce sectors, but it will only work for

successful fraud detection in e-commerce if accompanied by a careful retraining effort, a thoughtful priming of relevant features, and knowledge of each respective domain's behavior. A banking model that is predicated on significant and sudden high-value deviations will not easily transfer or marry itself to the inherently spiky and seasonal feature of e-commerce behaviors.

Moving forward, the outlook appears to be hybrid detection architectures, anonymized cross-platform intelligence sharing, and development of real-time behavioral analytics. This review recognizes several limitations. To begin with, the study substantially relies on technology platform case studies of major instances, which might not be representative of smaller operations with different profiles of fraud exposure and resource constraints. Secondly, the fraud techniques are continuously changing, and the performance benchmarks referred to may become obsolete as the adversaries change. The research on adversarial machine learning reveals that attackers deliberately test detection systems to find decision boundaries, and then they construct evasion attacks that exploit model vulnerabilities. This is an ongoing arms race that always requires model updating.[33] Thirdly, the regulatory frameworks that govern cross-sector data sharing are still different in various jurisdictions; thus, there are legal uncertainties regarding the proposed collaboration models. Lastly, this review mainly considers the technical literature, while it also acknowledges that organizational factors, data governance practices, analyst expertise, and risk management culture have a great impact on the effectiveness of fraud detection in the real world, which is independent of the level of algorithmic sophistication.[34] Research to come should fill these gaps by conducting longitudinal studies on model performance degradation, empirical evaluation of federated learning implementation under realistic regulatory conditions, and comparative analyses of fraud detection capabilities in small versus large platforms. Technology cannot do it alone. The pace of regulatory change must keep up, and Regulators should reasonably interrogate the balance between enabling innovative, ethical, informative, and fair technologies. While inter-industry cooperation has historically been difficult, it may now become necessary due to fraud schemes taking advantage of poor coordination across industries.

It is time to take action. The research community will need to focus on developing privacy-protecting and standardized methods of data collaboration (e.g., federated learning methods). At the same time, industry consortia and regulators should coordinate to provide ethical guidelines and auditing standards to address algorithmic bias, allowing AI to protect us but not revisit these same injustices in predictive ways.

If these barriers can be removed, we might begin to think of a new class of AI-assisted systems that are not limited to a particular industry but can operate with similar accuracy across a range of settings. These systems would recognize fraud not only where it occurs, but also where it is likely to occur next. In a world where fraud is highly adaptive, adaptability itself may become the most important feature of any detection model.

## Acknowledgments

I would like to thank Samuel Lefcourt and Dr. Siddharth Krishnan for their guidance, insight, and constructive feedback, which helped improve the quality of this paper and strengthen its analysis. I attest that the ideas, graphics, and writing in this paper are entirely my own.

## References

1. Federal Trade Commission. New FTC Data Show Big Jump in Reported Losses to Fraud: $12.5 Billion in 2024 [Online]; Press Release, **2025**. https://www.ftc.gov/news-events/news/press-releases/2025/03/new-ftc-data-show-big-jump-report ed-losses-fraud-125-billion-2024.
2. Anderson, K. B. Consumer Fraud in the United States, 2011: The Third FTC Survey; Federal Trade Commission Staff Report, **2013**. https://www.ftc.gov/sites/default/files/documents/reports/consumer-fraud-united-states-2011-third-f tc-survey/130419fraudsurvey_0.pdf.
3. Kircanski, K.; Notthoff, N.; Shadel, D.; Vasquez, E.; Gonzales, E.; Isacoff, J.; Carstensen, L. L. Heightened Emotional States Increase Susceptibility to Fraud in Older Adults. J. Elder Abuse Negl. **2018**. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6202193/.
4. Xenoss. Real-Time AI Fraud Detection in Banking. https://xenoss.io/blog/real-time-ai-fraud-detection-in-banking.
5. IBM. AI Fraud Detection in Banking. https://www.ibm.com/think/topics/ai-fraud-detection-in-banking.
6. DataStax. Real-Time Fraud Detection for Financial Services. https://www.datastax.com/guides/real-time-fraud-dectection-for-financial-services.
7. Zhang, Y.; Liu, H. AI-Driven Fraud Detection in Digital Finance. Humanit. Soc. Sci. Commun. **2024**, 11, 3606. https://www.nature.com/articles/s41599-024-03606-0.pdf.
8. Devarakonda, R. R. Machine Learning Approach for Fraud Detection in a Financial Services Application. Int. J. Sci. Adv. Technol. **2023**, 1, 2878. https://www.ijsat.org/papers/2023/1/2878.pdf.
9. Ali, A.; Razak, S. A.; Othman, S. H.; Eisa, T. A.; Al-Dhaqm, A.; Nasser, M.; Elhassan, T.; Elshafie, H.; Saif, A. Financial Fraud Detection Based on Machine Learning: A Systematic Literature Review. Appl. Sci. **2022**, 12 (19), 9637. https://www.mdpi.com/2076-3417/12/19/9637.
10. Chen, Y.; Zhao, C.; Xu, Y.; Nie, C. Year-Over-Year Developments in Financial Fraud Detection via Deep Learning: A Systematic Literature Review. Preprint posted on arXiv [Online]. **2025**. arXiv:2502.00201v1.
11. AIMultiple. Fake Review Detection: How AI Helps Identify Fraudulent Reviews. https://research.aimultiple.com/fake-review-detection/.
12. Academy of Accounting and Financial Studies Journal. Role of Artificial Intelligence in Financial Fraud Detection. https://www.abacademies.org/articles/role-of-artificial-intelligence-in-financial-fraud-detection.pdf.
13. GSC Online Press. AI-Based Fraud Detection in Financial Services. GSC Adv. Res. Rev. **2025**, 25 (4). https://gsconlinepress.com/journals/gscarr/sites/default/files/GSCARR-2025-0024.pdf.
14. Carcillo, F.; Dal Pozzolo, A.; Le Borgne, Y.-A.; Caelen, O.; Bontempi, G.; Mazzer, Y. Combining Unsupervised and Supervised Learning in Credit Card Fraud Detection. Preprint posted on arXiv [Online]. **2019**. arXiv:1903.04687.
15. Dal Pozzolo, A.; Boracchi, G.; Caelen, O.; Alippi, C.; Bontempi, G. Credit Card Fraud Detection: Realistic Modeling and Novel Learning Strategy. IEEE Trans. Neural Netw. Learn. Syst. **2017**, 29 (8), 3784–3797. https://ieeexplore.ieee.org/document/7553459.


DOI: 10.36838/IJHSR89.237

16. Bhattacharyya, S.; Jha, S.; Tharakunnel, K.; Westland, J. C. Data Mining for Credit Card Fraud: A Comparative Study. Decis. Support Syst. **2011**, 50 (3), 602–613. https://euro.ecom.cmu.edu/resources/elibrary/epay/1-s2.0-S0167923610001326-main.pdf.
17. Fiore, U.; De Santis, A.; Perla, F.; Zanetti, P.; Palmieri, F. Using Generative Adversarial Networks for Improving Classification Effectiveness in Credit Card Fraud Detection. Expert Syst. Appl. **2019**, 98, 420–430. https://www.iris.unisa.it/retrieve/e2915b35-9847-8981-e053-6605fe0a83a3/ins_Alfredo.pdf .
18. Chen, C.; Li, X.; Huang, J. Autoencoder-Based Anomaly Detection for Financial Fraud. Soft Comput. **2023**, 27, 12345–12360. https://e-jurnal.rokania.ac.id/index.php/jictas/article/view/431/290
19. Wang, X.; He, X.; Gao, W.; Yang, F.; An, Z.; Chen, J. How Effective Are Graph Neural Networks in Fraud Detection for Network Data? Preprint posted on arXiv [Online]. **2021**. arXiv:2105.14361.
20. Mir, A. Q.; Khan, F. Y.; Chishti, M. A. Online Fake Review Detection Using Supervised Machine Learning and BERT Model. Preprint posted on arXiv [Online]. **2023**. arXiv:2301.03225.
21. Gao, G.; Burtch, G.; Greenwood, B. N. Digitization and Fraud: Evidence from a Natural Experiment. Manag. Sci. **2022**. https://pubsonline.informs.org/doi/10.1287/mnsc.2021.4188.
22. Leung, E.; Paolacci, G.; Puntoni, S. Man versus Machine: Resisting Automation in Identity-Based Consumer Behavior. J. Mark. Res. **2018**. https://journals.sagepub.com/doi/10.1177/0022243718818423.
23. Juniper Research. AI Financial Fraud Detection Infographic. https://www.juniperresearch.com/resources/infographics/ai-financial-fraud-detection-infographic
24. Chua, C. E. H.; Wareham, J. Fighting Internet Auction Fraud: An Assessment and Proposal. Computer **2004**, 37 (10), 31–37. https://ieeexplore.ieee.org/document/1341833.
25. Kauffman, R. J.; Wood, C. A. Doing Their Bidding: An Empirical Examination of Factors That Affect a Buyer's Utility in Internet Auctions. Electron. Mark. **2006**. https://link.springer.com/article/10.1007/s10799-006-9181-4
26. Humphreys, A.; Latour, K. A. Framing the Game: Assessing the Impact of Cultural Representations on Consumer Perceptions of Legitimacy. J. Consum. Res. **2013**, 40 (4), 773–792. https://academic.oup.com/jcr/article/40/4/773/1798854.
27. Bolton, R. N.; Kannan, P. K.; Bramlett, M. D. Implications of Loyalty Program Membership and Service Experiences for Customer Retention and Value. J. Acad. Mark. Sci. **2000**, 28 (1), 95–108. https://link.springer.com/article/10.1177/0092070300281009.
28. Einav, L.; Klenow, P. J.; Klopack, B.; Levin, J. D.; Levin, L.; Best, W. Assessing the Gains from E-Commerce. NBER Working Paper 25610, **2018**. https://www.nber.org/papers/w25610.
29. Brouthers, L. E.; Geisser, K. D.; Rothlauf, F. Explaining the Internationalization of iBusiness Firms. J. Int. Bus. Stud. **2016**. https://link.springer.com/article/10.1057/jibs.2015.20.
30. Amazon Web Services. How Mastercard Achieved Near-Zero Downtime Deployments for Fraud Detection. AWS Industries Blog [Online], **2023**. https://aws.amazon.com/blogs/industries/how-mastercard-achieved-near-zero-downtime-deploym ents-for-fraud-detection-2/.
31. Federal Trade Commission. How Scammers Try to Steal Your Life Savings [Infographic]. https://consumer.ftc.gov/articles/how-scammers-try-steal-your-life-savings-infographic.
32. Chopra, P.; Binwal, A. The Role of AI/ML in Enhancing Security and Fraud Detection in Digital Payments. Int. J. Fundam. Mod. Res. **2024**. https://www.ijfmr.com/papers/2024/6/30337.pdf.
33. Cartella, F.; Anunciacao, O.; Funabiki, Y.; Yamaguchi, D.; Akishita, T.; Elshocht, O. Adversarial Attacks for Tabular Data: Application to Fraud Detection and Imbalanced Data. Preprint posted on arXiv [Online]. **2021**. arXiv:2101.08030.
34. Ngai, E. W. T.; Hu, Y.; Wong, Y. H.; Chen, Y.; Sun, X. The Application of Data Mining Techniques in Financial Fraud Detection: A Classification Framework and Academic Review of Literature. Decis. Support Syst. **2011**, 51 (4), 1105–1114. https://www.sciencedirect.com/science/article/abs/pii/S0167923610001302.

## Author

Andy Wu is a high school senior who wants to pursue a major in applied mathematics in college. He has committed to the University of Florida, and in his free time, he likes playing basketball, watching all 4 major US sports leagues, and food volunteering.