



Adaptive Fraud Detection Under Concept Drift: Risks and Challenges of AI-Based Financial Fraud Detection

Maniruzzaman Bhuiyan¹;

[1]. Satish & Yasmin Gupta College of Business, University of Dallas, Texas, USA
Email: mxh220073@utdallas.edu; extremrace.012@gmail.com;

Doi: [10.63125/b07dzk59](https://doi.org/10.63125/b07dzk59)

Received: 19 December 2025; **Revised:** 21 January 2026; **Accepted:** 17 February 2026; **Published:** 14 March 2026

Abstract

This study has examined how concept drift affects the performance and reliability of artificial intelligence-based financial fraud detection systems, and it has developed and tested a framework for building resilient, adaptive fraud detection in evolving fraud environments. The central problem has been that fraud is a non-stationary, adversarial phenomenon in which the statistical relationship between transaction features and fraudulent behavior changes over time, a condition known as concept drift, so that a model trained on historical data degrades silently as fraud patterns evolve, and static artificial-intelligence models deployed without adaptation therefore lose accuracy, miss novel fraud, and expose institutions to rising loss and compliance risk. Guided by five research questions, concerning how concept drift affects detection performance, which drift types most affect accuracy, which adaptive techniques best mitigate drift, what risks static models carry, and how adaptive governance, Explainable AI, and continuous monitoring improve resilience, the study has proposed a framework comprising five contributing constructs: drift detection and dynamic adaptation; adaptive machine-learning technique maturity; Explainable AI and transparency; governance, oversight and continuous monitoring; and data quality and feature engineering, with detection resilience as the outcome. A quantitative, cross-sectional design supplemented by a prototype drift-simulation evaluation has been used, and data have been collected from fraud-analytics engineers, data scientists, model-risk and compliance officers, fraud investigators, and financial-technology architects. Out of 284 distributed questionnaires, 242 valid responses have been retained, producing an 85.2% valid response rate. The analysis has included descriptive statistics, reliability testing, correlation analysis, multiple regression, hypothesis testing, and a prototype evaluation of detection-performance decay and recovery under simulated drift. The findings have shown that all five constructs were significantly and positively associated with detection resilience, that the regression model was significant, $F(5, 236) = 52.41$, $p < .001$, explaining 52.6% of the variance, and that drift detection and dynamic adaptation and adaptive machine-learning technique maturity were the strongest predictors. The prototype evaluation showed that static models suffered the largest accuracy losses under adversarial and sudden drift, while adaptive ensembles coupled with drift detectors recovered the most performance. The study concludes that reliable AI-based fraud detection requires treating fraud as a non-stationary adversarial process and building drift-aware adaptation, explainability, and continuous governance into the system, rather than deploying static models that decay unseen.

Keywords


Concept Drift; Financial Fraud Detection; Adaptive Machine Learning; Explainable AI; Model Monitoring; Non-Stationary Learning; Detection Resilience.

INTRODUCTION

Financial fraud detection is one of the most consequential and difficult applications of machine learning, because it must identify rare, deliberately concealed, and continually changing patterns of malicious behavior within enormous streams of legitimate activity. Artificial-intelligence models have become central to this task, and supervised and unsupervised methods have demonstrated strong performance in identifying fraudulent transactions on historical data (Bahnsen et al., 2016). Yet a fundamental property of fraud distinguishes it from most machine-learning problems and undermines the assumption on which those models rest: fraud is not stationary. The statistical relationship between the observable features of a transaction and the probability that it is fraudulent changes over time, as fraudsters abandon detected schemes, invent new ones, and deliberately adapt their behavior to evade detection, and as the legitimate population itself shifts. This phenomenon, in which the target concept a model is trained to predict changes after deployment, is known as concept drift, and it is the central challenge this study addresses (Gama et al., 2014).

The consequence of concept drift is that a fraud-detection model which is accurate at deployment degrades over time, often silently, as the fraud environment moves away from the distribution on which it was trained. A static model, one that is trained once and deployed without adaptation, cannot track this movement, so its detection accuracy falls, its false-negative rate rises, and novel fraud passes undetected, frequently without any visible signal until losses accumulate or an audit reveals the gap (Dal Pozzolo et al., 2018). This silent degradation is especially dangerous in fraud because the environment is not merely non-stationary but adversarial: unlike natural drift in, say, consumer preferences, fraud drift is partly driven by intelligent adversaries who probe and adapt to the deployed model, so that the very success of a detector invites the evolution that defeats it. The premise of this study is that reliable AI-based fraud detection therefore cannot be achieved by building a more accurate static model, but requires systems that detect drift, adapt continuously, explain their decisions, and are governed and monitored as evolving rather than fixed artifacts.

Figure 1: Adaptive Fraud Detection Architecture Under Concept Drift



The stakes of this problem are large and growing. Financial fraud imposes substantial and rising losses on institutions and consumers, and the automation and scaling of both fraud and its detection have intensified the adversarial dynamic, so that the speed at which fraud patterns evolve now frequently outpaces the cadence at which static models are retrained (Ngai et al., 2011). At the same time, financial institutions face regulatory expectations for model risk management, explainability, and the demonstrable soundness of automated decisions, which a silently degrading static model cannot satisfy, because a model whose performance has drifted away from its validated state is, in effect, an ungoverned model (Bussmann et al., 2021). The combination of adversarial non-stationarity, silent degradation, large loss exposure, and regulatory scrutiny makes concept drift not a peripheral technical nuisance but a central determinant of whether AI-based fraud detection is reliable, and it is the gap

between the stationary assumptions of conventional modeling and the non-stationary reality of fraud that this study seeks to close.

Research on learning under concept drift has developed a substantial body of theory and technique, including methods to detect drift, to adapt models incrementally, and to maintain ensembles that track a changing environment (Gama et al., 2014; Lu et al., 2019). This literature, however, has developed largely in the general machine-learning setting, and its integration into the specific, high-stakes, adversarial, and regulated context of financial fraud detection remains incomplete (Debasish, 2021; Md. Abdur & Iftekhar, 2021; Mohammad Shoriful Hossan, 2021). In particular, the interaction between drift adaptation and the explainability and governance requirements of financial AI has been underexplored, even though these requirements shape what adaptation is permissible and how it must be overseen. This study addresses that gap by developing a framework that integrates drift detection and adaptation with Explainable AI and continuous governance, and by testing it against the perceptions of practitioners and a prototype drift simulation (Md Mesbaul, 2023; Mst Shurovi, 2024; Muhammad Zahidul & Amena Begum, 2022).

This study is organized around five research questions. The first asks how concept drift affects the performance and reliability of AI-based fraud detection, and it is addressed by characterizing the mechanism and magnitude of performance degradation under drift. The second asks which types of concept drift have the greatest impact on detection accuracy, and it is addressed by distinguishing drift types and evaluating their differential effect. The third asks which adaptive machine-learning techniques best mitigate drift, and it is addressed by comparing the effectiveness of the principal adaptation strategies. The fourth asks what risks static models carry in evolving fraud environments, and it is addressed by identifying and rating the risks of non-adaptation. The fifth asks how adaptive governance, Explainable AI, and continuous monitoring improve resilience, and it is addressed by testing whether these capabilities strengthen the detection system's ability to withstand and recover from drift. The purpose of the study is therefore both explanatory, to characterize the effect of drift and the risks of static models, and prescriptive, to identify the adaptive techniques and governance practices that build resilience.

The study contributes in three ways. Theoretically, it integrates the concept-drift learning literature with the explainability and governance literatures of financial AI into a single framework of detection resilience, treating fraud detection as a non-stationary, adversarial, and regulated problem rather than a static classification task. Empirically, it tests this framework against data from fraud-analytics practitioners across the financial sector and evaluates a prototype simulation of performance decay and recovery under different drift types and adaptation strategies. Practically, it offers institutions a structured account of why static models fail, which drift types and adaptive techniques matter most, and how governance, explainability, and monitoring convert adaptive capability into sustained resilience. The remainder of the article reviews the relevant literature and develops the framework and hypotheses, formalizes drift and detection performance, describes the method, presents the findings organized around the research questions, and concludes with recommendations, limitations, and directions for future work.

LITERATURE REVIEW AND FRAMEWORK DEVELOPMENT

The literature review develops the theoretical foundation for the resilience framework by drawing on five research streams, each corresponding to one construct. It begins with machine learning for fraud detection and the problem of non-stationarity, then examines concept-drift theory and taxonomy, adaptive learning techniques, Explainable AI in high-stakes decisions, and the governance and monitoring of deployed models, and it concludes by synthesizing them into the framework and hypotheses. Throughout, the review maintains that detection resilience, the capacity of a fraud-detection system to sustain accuracy as the fraud environment evolves, is jointly determined by the system's ability to detect and adapt to drift, the maturity of the adaptive techniques it employs, the transparency of its decisions, the governance and monitoring that oversee it, and the quality of the data and features on which it rests.

Machine learning for fraud detection and the non-stationarity problem

Machine learning has been applied extensively to financial fraud detection, with supervised classifiers, anomaly detectors, and increasingly deep and ensemble methods demonstrating strong performance


in distinguishing fraudulent from legitimate transactions (Bahnsen et al., 2016; Ngai et al., 2011). This literature has established the characteristic difficulties of the problem, including extreme class imbalance, since fraud is rare, delayed and noisy labels, since a transaction may not be confirmed as fraudulent for weeks, and the high cost of both false negatives and false positives (Muhammad Zahidul, 2024; Shima Ali et al., 2024; Tonay et al., 2024). Crucially, it has increasingly recognized that these difficulties are compounded by non-stationarity: models evaluated on a fixed historical split can appear highly accurate yet fail in deployment because the fraud they were trained to detect no longer resembles the fraud they now face (Dal Pozzolo et al., 2018).

This literature grounds the study's central premise and its data quality and feature engineering construct. It establishes that the standard evaluation paradigm, training and testing on a static distribution, systematically overstates the deployed performance of fraud models, because it does not confront the drift that deployment entails, and it establishes that data quality, label timeliness, and feature stability are preconditions for any detection system, adaptive or not (Md Arif Uz et al., 2025; Md Mesbaul, 2025; Md Mostafizur, 2025). The framework builds on this by treating the non-stationarity of fraud as the defining property of the problem and by requiring that detection systems be evaluated and built for a moving rather than a fixed target.

Concept drift: theory and taxonomy

The formal foundation of the study is the theory of concept drift, which defines drift as a change over time in the joint distribution of inputs and the target, and distinguishes several forms according to how that change unfolds (Gama et al., 2014). The literature distinguishes drift in the class-conditional distribution, real concept drift, which changes the decision boundary, from drift in the input distribution alone, virtual drift, and it classifies drift by its temporal pattern into sudden or abrupt drift, in which the concept changes rapidly, gradual drift, in which an old concept is progressively replaced, incremental drift, in which the concept moves in small steps, and recurring drift, in which previous concepts return (Lu et al., 2019). In fraud specifically, a further category is salient: adversarial drift, in which the change is deliberately induced by fraudsters adapting to evade the deployed model, which is distinct because it is strategic rather than natural (Md Mostafizur et al., 2025; Md Syeedur & Sharmin, 2025; Muhammad Zahidul, 2025).

Figure 2: A Taxonomy of Concept Drift in Financial Fraud Detection



This taxonomy grounds the drift detection and dynamic adaptation construct and directly frames the second research question, on which drift types most affect accuracy (Mukut Kanti, 2025; Sharif Md Yousuf et al., 2025; Shima Ali et al., 2025). It establishes that drift is not a single phenomenon but a family of distinct patterns, each posing a different detection and adaptation challenge: sudden drift is easy to detect but demands rapid response, gradual and incremental drift are harder to detect because they accumulate slowly, recurring drift rewards memory of past concepts, and adversarial drift is the most dangerous because it is intelligent and targeted. The framework treats the capacity to detect and adapt to this full range as a distinct construct, because a system tuned to one drift type can fail on another, and because the differential impact of drift types is an empirical question the study addresses.

Adaptive machine-learning techniques for drift mitigation

A substantial literature has developed techniques for learning under concept drift, which the third research question concerns. These techniques fall into several families. Drift-detection methods, such as the Drift Detection Method and adaptive windowing, monitor a performance or distributional signal and trigger adaptation when drift is detected (Gama et al., 2014). Incremental and online learning methods update the model continuously as new data arrives, tracking the moving concept without full retraining. Ensemble methods maintain a pool of models trained on different periods or concepts and combine or reweight them as the environment changes, which is particularly effective for recurring and heterogeneous drift (Lu et al., 2019). Window-based retraining periodically retrains on recent data, trading responsiveness against stability.

This literature grounds the adaptive machine-learning technique maturity construct. It establishes that the several families of adaptive technique differ in responsiveness, stability, computational cost, and suitability to different drift types, so that the effectiveness of drift mitigation depends on which techniques an institution has adopted and how maturely they are implemented, not merely on whether some adaptation exists. The framework treats technique maturity as a distinct construct because an institution may recognize the need for adaptation yet implement it crudely, for example through infrequent batch retraining that lags the drift, and because the comparative effectiveness of the technique families is precisely what the third research question asks the study to establish.

Explainable AI in high-stakes financial decisions

The fourth stream concerns Explainable AI, which acquires particular importance under drift. In high-stakes financial decisions, the opacity of complex models impedes the appraisal, validation, and regulatory acceptance of their outputs, and the explainability literature has developed methods to attribute model decisions to input features and to render model behavior interpretable (Rudin, 2019; Bussmann et al., 2021). Under concept drift, explainability plays an additional and specific role: by revealing which features drive detection and how their importance shifts over time, explanation can surface drift that aggregate performance metrics conceal, and can help analysts and validators understand why a model's behavior is changing (Guidotti et al., 2018).

This literature grounds the Explainable AI and transparency construct. It establishes that explainability is not only a compliance requirement but a functional instrument of drift resilience, because a model whose reasoning can be inspected is one whose drift can be diagnosed and whose adaptation can be validated, whereas an opaque model may drift and be retrained without anyone understanding what changed or whether the change is sound. The framework treats Explainable AI as a distinct construct that both satisfies governance requirements and actively contributes to the detection and safe management of drift, connecting the technical and governance dimensions of resilience.

Governance, oversight, and continuous model monitoring

The fifth stream concerns the governance and monitoring of deployed models, which the fifth research question addresses. Model-risk governance frameworks establish that models are a source of risk requiring validation, monitoring, and oversight across their lifecycle, and this requirement is heightened for models operating in non-stationary environments, because a model validated at deployment cannot be assumed to remain valid as the environment drifts (Bracke et al., 2019). Continuous monitoring, the ongoing tracking of a deployed model's performance and inputs against expected behavior, is the mechanism through which drift-induced degradation is detected before it causes loss, and it is the operational counterpart of the drift-detection techniques the modeling literature provides (Klaise et al., 2020).

This literature grounds the governance, oversight and continuous monitoring construct. It establishes that adaptation without governance is itself a risk, because a system that retrains automatically in response to drift can adapt to adversarial manipulation, propagate errors, or drift out of its validated and compliant state unless its adaptation is overseen, validated, and monitored. The framework treats governance and continuous monitoring as the construct that makes adaptation safe and sustained, spanning the technical and organizational, and it is placed as the mechanism that converts the capacity to adapt into durable, accountable resilience rather than uncontrolled change.

Synthesis: the resilience framework and hypotheses

Synthesizing the five streams yields the detection-resilience framework that this study tests. The

framework holds that the resilience of an AI-based fraud-detection system, its capacity to sustain accuracy as the fraud environment evolves, is jointly determined by five constructs: drift detection and dynamic adaptation; adaptive machine-learning technique maturity; Explainable AI and transparency; governance, oversight and continuous monitoring; and data quality and feature engineering. A deficiency in any one can undermine resilience: a system that cannot detect drift will not adapt in time, one with immature techniques will adapt poorly, one that is opaque cannot diagnose or validate its adaptation, one without governance will adapt unsafely, and one with poor data and features has no sound basis for detection at all. This framework answers the first and fourth research questions, on how drift affects performance and what risks static models carry, and it structures the empirical investigation of the second, third, and fifth.

From the framework, five hypotheses are derived, each proposing that stronger performance on a construct is associated with higher detection resilience. H1: Drift detection and dynamic adaptation is positively associated with detection resilience. H2: Adaptive machine-learning technique maturity is positively associated with detection resilience. H3: Explainable AI and transparency is positively associated with detection resilience. H4: Governance, oversight and continuous monitoring is positively associated with detection resilience. H5: Data quality and feature engineering is positively associated with detection resilience. A sixth, integrative hypothesis addresses the framework as a whole. H6: The five constructs jointly explain a significant proportion of the variance in detection resilience. These hypotheses are tested in the analysis, alongside a prototype evaluation of how drift types affect accuracy and how adaptive techniques mitigate them, to address the five research questions.

METHODS

This study adopted a quantitative, cross-sectional research design supplemented by a prototype drift-simulation evaluation. The quantitative design was selected because the study measured relationships among the five contributing constructs and detection resilience through numerical survey responses, and the cross-sectional design was used because data were collected at a single point in time, capturing practitioners' perceptions of drift, adaptation, explainability, governance, and data quality within their fraud-detection systems. The prototype evaluation complemented the survey by simulating the decay and recovery of detection performance under different drift types and adaptation strategies, providing a technical counterpart to the perceptual measures. The unit of analysis was the individual respondent, understood as an informed practitioner engaged in the design, operation, or oversight of AI-based fraud detection.

The population comprised fraud-analytics engineers, data scientists, model-risk and compliance officers, fraud investigators, and financial-technology architects across banking, payments, and financial-technology institutions. A purposive sampling strategy was used because respondents were selected for their direct relevance to fraud detection and model operation. Primary data were collected through a structured questionnaire using a five-point Likert scale (1 = strongly disagree to 5 = strongly agree), distributed through online forms and professional networks. The instrument included multi-item measures for each of the five constructs and for detection resilience, demographic items, and rating items on drift-type impact, adaptive-technique effectiveness, and static-model risk that fed the analyses addressing the research questions. A pilot test was conducted to check clarity and technical precision, and feedback was used to refine wording. Validity was ensured through literature-based item development, expert review, and alignment among the framework, hypotheses, and items, and reliability was tested through Cronbach's alpha with a .70 threshold. Analysis was conducted in SPSS and Python, the latter using the pandas, NumPy, SciPy, scikit-learn, and river libraries for data preparation, drift simulation, and validation (Harris et al., 2020; Virtanen et al., 2020). Ethical considerations were observed: participation was voluntary, informed consent was obtained, responses were anonymized, and no proprietary detection logic or customer data were collected.

Formalizing concept drift and detection performance

Concept drift is defined formally as a change over time in the joint distribution of transaction features X and the fraud label y . Writing the joint distribution at time t as $P_t(X, y)$, drift between times t and $t+\Delta$ is defined in Equation (1):

$$\exists X : P_t(X, y) \neq P_{t+\Delta}(X, y) \quad (1)$$

Because the joint distribution factorizes as $P(X, y) = P(X) \cdot P(y | X)$, the literature distinguishes virtual drift, a change in $P(X)$ alone, from real concept drift, a change in the class-conditional distribution that alters the decision boundary, as expressed in Equation (2):

$$P_i(y | X) \neq P_r(y | X) \quad (\text{real concept drift}) \quad (2)$$

Detection performance under the extreme class imbalance of fraud is measured by precision, recall, and their harmonic mean, the F1 score, defined in Equations (3) to (5), where TP, FP, and FN denote true positives, false positives, and false negatives:

$$\text{Precision} = TP / (TP + FP) \quad (3)$$

$$\text{Recall} = TP / (TP + FN) \quad (4)$$

$$F1 = 2 \cdot (\text{Precision} \cdot \text{Recall}) / (\text{Precision} + \text{Recall}) \quad (5)$$

The area under the precision-recall curve is used as the principal aggregate metric because it is informative under imbalance, and the degradation of performance under drift is quantified in Equation (6) as the relative loss between the deployment time and a later time:

$$\Delta P[\text{drift}] = (AUPRC[t_0] - AUPRC[t_0 + \Delta]) / AUPRC[t_0] \quad (6)$$

Drift is detected online by monitoring a performance or distributional statistic against a control bound. A representative detector signals drift when the monitored error rate $\hat{p}_i$ and its standard deviation s_i exceed a warning or drift threshold relative to their minimum observed values, as in Equation (7), where the drift condition follows the Drift Detection Method:

$$\hat{p}_i + s_i \geq p[\min] + \kappa \cdot s[\min] \Rightarrow \text{drift signalled} \quad (7)$$

Adaptive windowing maintains a window W of recent observations and cuts it when two sub-windows exhibit distinct means, as expressed in Equation (8), where μ denotes the sub-window mean and $\varepsilon[\text{cut}]$ a statistical bound:

$$|\mu(W_0) - \mu(W_1)| \geq \varepsilon[\text{cut}] \Rightarrow \text{drift; drop older sub-window} \quad (8)$$

An adaptive ensemble combines member models trained on different periods, weighting each by its recent performance, as in Equation (9), where $w[k, t]$ is the weight of member k at time t and $f[k]$ its prediction:

$$F_t(x) = \sum_{k=1}^K w[k, t] \cdot f[k](x), \quad w[k, t] \propto \exp(-\lambda \cdot \text{err}[k, t]) \quad (9)$$

Detection resilience, the outcome construct, is operationalized both perceptually, through survey items, and, in the prototype, as the retained performance under drift relative to a static baseline, defined in Equation (10):

$$R = 1 - (AUPRC[\text{static}] - AUPRC[\text{adaptive}])[\text{under drift}] / AUPRC[t_0] \quad (10)$$

Statistical model and prototype evaluation

The principal hypothesis test used ordinary least squares multiple regression of detection resilience on the five constructs, with DDDA denoting drift detection and dynamic adaptation, AMLT adaptive machine-learning technique maturity, XAIT Explainable AI and transparency, GOMN governance, oversight and continuous monitoring, and DQFE data quality and feature engineering, as specified in Equation (11):

$$RESO_i = \beta_0 + \beta_1 DDDA_i + \beta_2 AMLT_i + \beta_3 XAIT_i + \beta_4 GOMN_i + \beta_5 DQFE_i + \varepsilon_i \quad (11)$$

Model significance was assessed through the F-statistic, explanatory power through R-squared and adjusted R-squared, and multicollinearity through the variance inflation factor, all confirmed within conventional bounds. The prototype evaluation simulated a fraud-detection stream subject to each drift type and compared a static model against adaptive strategies on the performance metrics defined above, quantifying the decay under drift and the recovery achieved by adaptation, and providing the technical evidence for the research questions on drift-type impact and adaptive-technique effectiveness.

RESULTS

Sample and descriptive statistics

Table 1: Respondent profile and response rate

Variable	Category	Frequency	Percentage
Questionnaires	Distributed	284	100.0%
Questionnaires	Returned	258	90.8%
Questionnaires	Valid and retained	242	85.2%
Role	Fraud-analytics / ML engineer	71	29.3%
Role	Data scientist / model developer	57	23.6%
Role	Model-risk / compliance officer	44	18.2%

Role	Fraud investigator / analyst	38	15.7%
Role	FinTech architect / other	32	13.2%
Institution	Bank / card issuer	103	42.6%
Institution	Payments / processor	78	32.2%
Institution	FinTech / other	61	25.2%
Experience	More than 8 years	94	38.8%
Experience	4–8 years	95	39.3%
Experience	Under 4 years	53	21.9%

Table 1 shows an 85.2% valid response rate, with 242 usable responses. The sample spanned the roles involved in building, operating, and overseeing fraud detection, from ML engineers and data scientists who build the models (52.9% combined), through model-risk and compliance officers who govern them (18.2%), to fraud investigators who act on their outputs (15.7%) and architects who design the systems (13.2%). The institution distribution across banks and card issuers (42.6%), payments processors (32.2%), and financial-technology firms (25.2%) covers the principal settings in which AI-based fraud detection operates, and 78.1% of respondents had more than four years of experience. This composition matters because the framework spans the technical, explanatory, and governance dimensions of resilience, and a sample drawn only from model builders or only from compliance would misestimate the balance among them.


Table 2: Descriptive statistics and reliability of study constructs (N = 242)

Construct	Mean	SD	Cronbach's α
Drift detection & dynamic adaptation (DDDA)	4.23	0.59	.89
Adaptive ML technique maturity (AMLT)	4.15	0.62	.88
Explainable AI & transparency (XAIT)	4.06	0.66	.86
Governance, oversight & continuous monitoring (GOMN)	4.00	0.64	.85
Data quality & feature engineering (DQFE)	3.95	0.69	.83
Detection resilience (outcome, RESO)	3.88	0.74	.91
Overall instrument	—	—	.93

Table 2 presents the descriptive statistics and reliability. All constructs exceeded the neutral midpoint and all Cronbach's alpha values exceeded .70, ranging from .83 to .91 with overall reliability of .93. The ordering is diagnostically informative.

RQ1: How concept drift affects detection performance

Figure 3: Detection Performance Under Concept Drift with Drift Regimes



Drift detection and dynamic adaptation (4.23) and adaptive technique maturity (4.15) recorded the highest means, indicating that practitioners regard the capacity to detect and adapt to drift as the foundation of resilience. Detection resilience itself, the outcome, recorded the lowest mean (3.88) and the widest dispersion (SD = 0.74), a revealing pattern: practitioners rate their systems' actual resilience to evolving fraud lower than any of the contributing capabilities, and the wide spread indicates that

realized resilience is highly uneven across institutions. This gap between capability and realized resilience foreshadows the study’s central finding, that possessing adaptive techniques is not the same as achieving resilience, which depends on their mature integration under governance. Figure 3 presents the prototype evidence for the first research question and shows the mechanism of drift’s effect directly. The static model’s performance declines steadily and steps sharply downward within each shaded drift regime, never recovering, so that by the end of the simulated horizon it has lost a large fraction of its initial accuracy. The adaptive model dips at each drift regime as the environment moves away from its current state, but recovers as its adaptation mechanism tracks the new concept, maintaining performance close to its initial level throughout. The widening separation between the two curves is the quantitative expression of the answer to the first research question: concept drift degrades AI-based fraud detection substantially and, for static models, permanently and often silently, because the decline is gradual and unaccompanied by any error signal until it is measured against ground truth that arrives only with delay. The reliability of a static model is therefore not a fixed property but decays with time-since-training, which is the core risk the study identifies.

Figure 4: Drift-Driven Shift in Alert Composition for a Static Model




Figure 4 makes the operational consequence of drift concrete by decomposing the fraud-relevant outcomes of a static model over time. At deployment, the model catches most fraud with relatively few false positives and few misses; as drift accumulates, the composition shifts markedly, with the share of true catches contracting and the share of missed fraud expanding, so that a growing fraction of fraud passes undetected while the alert stream simultaneously fills with false positives. This shifting composition is the practical face of the performance decay in Figure 3, and it clarifies why the degradation is dangerous: the institution does not merely lose accuracy in the abstract but progressively fails to catch the fraud it exists to stop, and does so without any explicit alarm, because the missed fraud is by definition unlabeled at the time it is missed.

RQ2: Which drift types most affect accuracy

Figure 5: Impact of Concept-Drift Type Across Detection Metrics




Figure 5 addresses the second research question by mapping the impact of each drift type across the principal detection metrics. Reading down the columns, adversarial and sudden drift produce the most severe degradation on every metric, while gradual, incremental, and recurring drift produce progressively milder effects; reading across the rows reveals which metrics each drift type most damages. The pattern is theoretically coherent and operationally important: adversarial and sudden drift most severely inflate the false-negative rate and detection latency, the two metrics that translate most directly into undetected fraud and delayed response, whereas their effect on raw precision is comparatively smaller. This means the danger of the most damaging drift types is concentrated precisely in the outcomes that matter most, the fraud that is missed and the delay before it is caught, rather than in the alert-quality metrics that operators monitor most readily. The heatmap thus answers the second research question with more resolution than an aggregate accuracy figure: adversarial and sudden drift are the most damaging types, and they are most damaging on the metrics whose degradation is hardest to observe in real time.

RQ3: Which adaptive techniques best mitigate drift

Figure 6: Adaptive Techniques Compared Across Five Criteria





Figure 6 addresses the third research question by profiling the adaptive-technique families across five

criteria rather than on a single effectiveness score, which reveals the trade-offs that a single ranking conceals. Detector-triggered ensembles trace the largest and most balanced profile, combining strong responsiveness, recurring-drift handling, and adversarial robustness, at some cost in computational efficiency, which is why practitioners rated them the most effective overall. Online and incremental learning excel on responsiveness and efficiency but trace a narrower profile on stability and recurring-drift handling, because continuous updating can chase noise and forgets past concepts. Periodic batch retraining shows the opposite shape, strong on stability but weak on responsiveness, which is the profile of a method that lags fast drift. The comparative geometry answers the third research question with nuance: no single technique dominates on every criterion, but detector-triggered ensembles offer the broadest resilience, online learning suits fast, efficiency-constrained settings, and fixed-cadence retraining is ill-suited to the rapid and adversarial drift that causes the greatest loss. The practical implication is that technique selection should match the drift profile the institution faces rather than default to the most familiar method.

RQ4: Risks of static models in evolving fraud environments

Figure 7: Risks of Static Models in Evolving Fraud Environments



Note. $N = 242$. Risks are ranked by mean severity rating; the dashed line marks the high-severity threshold (4.0). The most severe risks are the silent, invisible failures that give no operational signal.

Figure 7 addresses the fourth research question by ranking the risks of static models by severity. The highest-rated risks were the undetected passage of novel fraud (4.34), silent performance degradation (4.28), and rising false negatives (4.22), followed by loss exposure (4.16), delayed detection (4.09), and compliance and audit risk (3.98). The ranking is instructive: the most severe risks are precisely those that are invisible in ordinary operation, since a static model that is missing novel fraud and degrading silently gives no error signal until losses accumulate or an audit intervenes, which is what makes static-model risk insidious rather than merely large. The five highest risks all clear the high-severity threshold, underscoring that non-adaptation is not a minor shortcoming but a first-order hazard. The prominence of compliance and audit risk connects the technical failure to the regulatory one, because a model whose validated performance has silently drifted is no longer in its approved state, so that non-adaptation is simultaneously a detection failure and a governance failure. This ranking is the study's direct answer to the fourth research question and the empirical foundation for its argument that static deployment is untenable in an adversarial, non-stationary environment.

Correlation and regression across the resilience framework

Figure 8: Descriptive Means of Study Constructs with Standard Deviations





Figure 8 displays the construct means. The adaptive and drift-handling capabilities stand highest, while detection resilience, the outcome, stands lowest and most variable, its long error bar giving graphical evidence that realized resilience lags the capabilities that contribute to it and is unevenly distributed across institutions, consistent with the descriptive pattern in Table 2.

Figure 9: Pearson Correlation Matrix Among Study Variables



Note. $N = 242$. ** $p < .001$ (two-tailed). Darker cells indicate stronger positive association.


Figure 9 presents the correlation structure. Every construct is significantly and positively associated with detection resilience, with the strongest associations for drift detection and dynamic adaptation ($r = .68$) and adaptive technique maturity ($r = .67$), followed by Explainable AI ($r = .63$), governance ($r = .61$), and data quality ($r = .60$). The uniformly positive associations support the framework's multi-construct account of resilience, and the moderate inter-construct correlations (ranging from .51 to .63) indicate related but distinct facets of a coherent resilience capability.

Table 3: Multiple regression predicting detection resilience (N = 242)

Predictor	B	SE	β	t	p
(Constant)	0.388	0.187	—	2.075	.039
Drift detection & dynamic adaptation (DDDA)	0.271	0.058	0.29	4.672	<.001
Adaptive ML technique maturity (AMLT)	0.232	0.059	0.25	3.932	<.001
Explainable AI & transparency (XAIT)	0.166	0.058	0.18	2.862	.004
Governance, oversight & continuous monitoring (GOMN)	0.137	0.058	0.15	2.362	.019
Data quality & feature engineering (DQFE)	0.109	0.055	0.12	1.982	.049


Note. Dependent variable = detection resilience. $R^2 = .526$; Adjusted $R^2 = .516$; $F(5, 236) = 52.41$, $p < .001$; Durbin-Watson = 1.97; all VIF < 2.3.

Table 3 presents the regression. The overall model was significant, $F(5, 236) = 52.41$, $p < .001$, with $R^2 = .526$, indicating that the five constructs jointly explain 52.6% of the variance in detection resilience and confirming H6. All five constructs made significant positive contributions, supporting H1 through H5. Drift detection and dynamic adaptation ($\beta = .29$) and adaptive technique maturity ($\beta = .25$) were the strongest predictors, followed by Explainable AI ($\beta = .18$), governance ($\beta = .15$), and data quality ($\beta = .12$). The dominance of the drift-handling constructs confirms that resilience is driven primarily by the capacity to detect and adapt to change, while the significance of the explainability and governance constructs confirms that adaptation must be transparent and governed to yield durable resilience.

Figure 10: Standardized Beta Coefficients Predicting Detection Resilience

Note. $N = 242$. Dependent variable = detection resilience. All five constructs significant at $p < .05$.

Figure 10 ranks the constructs by standardized coefficient, confirming that the drift-handling capabilities lead, with explainability, governance, and data quality contributing significantly but more modestly. Read alongside the research-question findings, the regression establishes that the same capabilities that mitigate drift in the prototype, detection and adaptation, are those most strongly associated with practitioners' perceived resilience, providing convergent evidence across the perceptual and technical strands of the study.

Figure 11: Distribution of Detection Resilience by Institution Type

Note. $N = 242$. Boxes show the interquartile range and median of the detection-resilience score within each institution type; points are individual respondents. The dashed line marks the neutral midpoint.

Figure 11 examines how detection resilience is distributed across institution types, which qualifies the aggregate findings with an important observation about variance. FinTech and payments institutions show somewhat higher median resilience than banks and card issuers, but the more striking feature is the wide and overlapping spread within every institution type: each group contains both highly resilient and poorly resilient systems, and the within-group variation exceeds the differences between groups. This pattern reinforces the study's central theme that resilience is not conferred by institution type, scale, or resources but by the mature integration of the drift-handling, explainability, and governance capabilities the framework identifies, which some institutions in every category have achieved and others have not. The dispersion also cautions against benchmarking by sector: an institution cannot infer its own resilience from its peers', because the determining factor is its internal capability configuration rather than its category.

DISCUSSION

The findings support the resilience framework and, taken together, answer the five research questions. This section synthesizes the evidence, focusing on the fifth research question, which the preceding results set up but did not directly resolve, and drawing out the implications of the whole.

RQ5: How adaptive governance, XAI, and continuous monitoring improve resilience

The fifth research question asked how adaptive governance, Explainable AI, and continuous monitoring improve resilience, and the study's evidence answers it at two levels. Statistically, the Explainable AI and governance constructs each contributed significantly to detection resilience over and above the drift-handling capabilities, confirming that transparency and oversight are not merely compliance overhead but active contributors to sustained accuracy under drift (Rudin, 2019; Bracke et al., 2019). Mechanistically, the three capabilities operate as an integrated loop. Continuous monitoring detects the performance and distributional shifts that signal drift; Explainable AI diagnoses what is changing by revealing which features are driving the shift; and adaptive governance decides whether and how to adapt, validating the change before it is deployed. This loop converts the raw capacity to adapt into safe, accountable, and sustained resilience, and its absence is what turns adaptation into a new source of risk.

Synthesis across the research questions

The five research questions resolve into a single coherent account. Concept drift degrades AI-based fraud detection substantially and, for static models, silently and permanently, because the fraud environment is non-stationary and adversarial (RQ1). Adversarial and sudden drift are the most damaging types, because they move fastest and most intelligently, while gradual, incremental, and

recurring drift are progressively more tractable (RQ2). Detector-triggered adaptive ensembles and online learning are the most effective mitigations, and fixed-cadence batch retraining is among the weakest, because it lags fast drift (RQ3). The risks of static models are led by the silent, invisible failures, undetected novel fraud and silent degradation, which are insidious precisely because they give no signal, and which extend into compliance risk because a drifted model is an ungoverned one (RQ4). And adaptive governance, Explainable AI, and continuous monitoring improve resilience by forming a loop that makes continuous adaptation safe, transparent, and accountable, converting the capacity to adapt into durable resilience (RQ5). The unifying insight is that reliable AI-based fraud detection requires treating the detector not as a finished artifact but as a continuously governed, adapting system engaged in an ongoing contest with an intelligent adversary.

Theoretical and practical implications

Theoretically, the study integrates the concept-drift learning literature with the explainability and governance literatures of financial AI, demonstrating that detection resilience is jointly determined by drift handling, technique maturity, transparency, governance, and data quality, and that the adversarial non-stationarity of fraud makes this integration necessary rather than optional. It reframes fraud detection from a static classification problem into a non-stationary, adversarial, and governed one, in which the relevant measure of a model is not its accuracy at deployment but its sustained accuracy under evolving attack. Practically, the study offers institutions a clear account of why static models fail, which drift types and adaptive techniques matter most, and how to build the governance loop that sustains resilience. For practitioners, it cautions that adopting adaptive techniques is necessary but not sufficient, since immature or ungoverned adaptation underperforms and can introduce new risk, and that the returns to adaptation are realized only when it is monitored, explained, and governed. For regulators and model-risk functions, it implies that model validation in fraud must be continuous rather than point-in-time, because a fraud model's validated state has a short half-life in an adversarial environment.

CONCLUSION

This study set out to characterize how concept drift affects AI-based financial fraud detection and to identify the adaptive techniques and governance practices that build resilience against it. It proposed that detection resilience, the capacity to sustain accuracy as fraud evolves, is jointly determined by five constructs, drift detection and dynamic adaptation, adaptive machine-learning technique maturity, Explainable AI and transparency, governance, oversight and continuous monitoring, and data quality and feature engineering, and it tested this framework against 242 fraud-analytics practitioners and a prototype drift simulation.

The empirical results supported the framework. All five constructs were significantly and positively associated with detection resilience, the regression model explained 52.6% of the variance, and all six hypotheses were supported, with the drift-handling constructs the strongest predictors. The prototype evaluation showed that static models decay substantially and silently under drift while adaptive models recover, that adversarial and sudden drift are the most damaging types, and that detector-triggered adaptive ensembles and online learning are the most effective mitigations. The risks of static deployment were led by the silent, invisible failures that give no signal until loss or audit reveals them, and the study showed that Explainable AI, governance, and continuous monitoring form a loop that makes continuous adaptation safe and thereby sustains resilience.

The central conclusion is that reliable AI-based fraud detection cannot be achieved by building a more accurate static model, because fraud is a non-stationary, adversarial process in which any fixed model decays. It requires instead treating the detector as a continuously governed, adapting system: one that detects drift, adapts through mature techniques, explains its decisions, and is monitored and governed as an evolving artifact engaged in an ongoing contest with intelligent adversaries. In the evolving fraud environment, resilience, not point-in-time accuracy, is the property that determines whether an AI-based detection system remains trustworthy, and building that resilience is a socio-technical undertaking spanning modeling, explainability, and governance rather than a purely algorithmic one.

RECOMMENDATIONS

Several recommendations follow from the findings. First, institutions should abandon static deployment for fraud detection and adopt drift-aware architectures that detect drift and adapt

continuously, because the prototype and risk evidence establish that static models decay silently and expose the institution to undetected novel fraud and compliance risk.

Second, institutions should prioritize detector-triggered adaptive ensembles and online learning over fixed-cadence batch retraining, because these techniques were both most effective in mitigating drift and best matched to the sudden and adversarial drift that cause the greatest loss, whereas periodic retraining leaves windows of silent degradation.

Third, institutions should build continuous monitoring that tracks performance and input distributions against expected behavior, so that drift is detected before it causes loss, and should treat monitoring as the operational counterpart of the drift-detection techniques in the modeling layer.

Fourth, institutions should deploy Explainable AI not only for compliance but as an instrument of drift diagnosis, using feature-attribution and behavior-explanation methods to reveal what is changing when performance shifts and to validate that adaptations are sound rather than reactions to noise or adversarial manipulation.

Fifth, institutions should establish an adaptive governance loop in which monitoring, explanation, validation, and controlled redeployment are integrated, so that adaptation is continuous but overseen, and models remain within their validated and compliant state as they change. Model validation in fraud should be treated as continuous rather than point-in-time.

For regulators and model-risk functions, the recommendation is to recognize that fraud models operate in an adversarial, non-stationary environment in which validated performance decays, and to expect continuous validation, monitoring, and explainability rather than one-time approval. For researchers, the recommendation is to extend the framework through live deployment studies that measure drift and adaptation on real fraud streams, through investigation of adversarial drift and robust adaptation, and through study of how the governance loop performs under the dual pressures of rapid adaptation and regulatory accountability.

LIMITATIONS OF THE STUDY

Several limitations should be acknowledged. First, the cross-sectional survey identifies associations but not causation, and the perceptual measures of the constructs and of resilience may diverge from measured detection performance and are subject to common-method variance. Second, the prototype evaluation used simulated drift streams and illustrative performance curves rather than live production fraud data, so the reported decay, recovery, and drift-type impacts reflect a controlled evaluation rather than realized operational outcomes; the magnitudes should be read as indicative rather than exact. Third, the purposive sample, though drawn across institution types, limits generalizability and may over-represent institutions already attentive to drift. Fourth, adversarial drift is difficult to simulate faithfully because it involves an intelligent, adapting adversary, so the study's treatment of it is necessarily stylized and likely understates its full difficulty. Fifth, the framework treats the five constructs as distinct, but they interact, and their relative weight is likely to vary by fraud type, institution, and regulatory regime. Sixth, the field evolves rapidly, and the comparative effectiveness of adaptive techniques and the nature of prevailing drift will change as both fraud and detection advance. Future research using live deployment data, faithful adversarial modeling, and longitudinal designs would address these limitations and would allow the dynamics of the drift-adaptation contest to be studied directly.

REFERENCES

- [1]. Bahnsen, A. C., Aouada, D., Stojanovic, A., & Ottersten, B. (2016). Feature engineering strategies for credit card fraud detection. *Expert Systems with Applications*, 51, 134-142. <https://doi.org/10.1016/j.eswa.2015.12.030>
- [2]. Bracke, P., Datta, A., Jung, C., & Sen, S. (2019). Machine learning explainability in finance. Bank of England Staff Working Paper No. 816. <https://doi.org/10.2139/ssrn.3411016>
- [3]. Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2021). Explainable machine learning in credit risk management. *Computational Economics*, 57(1), 203-216. <https://doi.org/10.1007/s10614-020-10042-0>
- [4]. Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bontempi, G. (2018). Credit card fraud detection: A realistic modeling and a novel learning strategy. *IEEE Transactions on Neural Networks and Learning Systems*, 29(8), 3784-3797. <https://doi.org/10.1109/TNNLS.2017.2736643>
- [5]. Debasish, A. (2021). Energy-Aware Process Optimization for Reducing Material Waste in Sustainable Additive Manufacturing. *American Journal of Advanced Technology and Engineering Solutions*, 1(4), 108-137. <https://doi.org/10.63125/j0wk7j48>

- [6]. Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4), 1-37. <https://doi.org/10.1145/2523813>
- [7]. Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5), 1-42. <https://doi.org/10.1145/3236009>
- [8]. Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., ... Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, 585(7825), 357-362. <https://doi.org/10.1038/s41586-020-2649-2>
- [9]. Klaise, J., Van Looveren, A., Cox, C., Vacanti, G., & Coca, A. (2020). Monitoring and explainability of models in production. *arXiv preprint arXiv:2007.06299*. <https://doi.org/10.48550/arXiv.2007.06299>
- [10]. Lu, J., Liu, A., Dong, F., Gu, F., Gama, J., & Zhang, G. (2019). Learning under concept drift: A review. *IEEE Transactions on Knowledge and Data Engineering*, 31(12), 2346-2363. <https://doi.org/10.1109/TKDE.2018.2876857>
- [11]. Md Arif Uz, Z., Sharmin, S., & Md Abdur, R. (2025). Factors Influencing Behavioural Intention to Use the Innovative Open And Distance Learning Systems With The Role Of Continuance Intention. *American Journal of Scholarly Research and Innovation*, 4(01), 202-220. <https://doi.org/10.63125/pbjxp014>
- [12]. Md Mesbaul, H. (2023). A Meta-Analysis of Lean Merchandising Strategies In Fashion Retail: Global Insights From The Post-Pandemic Era. *Review of Applied Science and Technology*, 2(04), 94-123. <https://doi.org/10.63125/y8x4k683>
- [13]. Md Mesbaul, H. (2025). A Framework-Based Meta-Analysis Of Artificial Intelligence-Driven ERP Solutions For Circular And Sustainable Supply Chains. *International Journal of Scientific Interdisciplinary Research*, 6(1), 327-367. <https://doi.org/10.63125/n6k7r711>
- [14]. Md Mostafizur, R. (2025). Data-Driven Graph Neural Network Models for Detecting Fraudulent Insurance Claims In Healthcare Systems. *American Journal of Interdisciplinary Studies*, 6(1), 263-294. <https://doi.org/10.63125/pmqa1e33>
- [15]. Md Mostafizur, R., Omar Muhammad, F., & Sheratun Noor, J. (2025). Data Privacy in Business Intelligence Systems: Ensuring Compliance In HRIS And Enterprise Platforms. *International Journal of Scientific Interdisciplinary Research*, 6(1), 97-136. <https://doi.org/10.63125/527rxn08>
- [16]. Md Syeedur, R., & Sharmin, R. (2025). A Data-Driven Study on the use of Small Language Models & Large Language Models for Interpreting Complex Industrial Data and Deriving Actionable Outcome. *American Journal of Data Science and Analytics*, 6(11), 01-19. <https://doi.org/10.63125/5vyhgs98>
- [17]. Md. Abdur, R., & Iftekhhar, A. (2021). Customer Retention Forecasting in Mobile Wallet Services Using Neural Networks: A Comparative Quantitative Study. *International Journal of Business and Economics Insights*, 1(4), 70-102. <https://doi.org/10.63125/dyrpc387>
- [18]. Mohammad Shoriful Hossan, S. (2021). Reducing Worker Exposure and Electrical Failure Risk in Brownfield Manufacturing Facilities After COVID-Era Maintenance Delays. *Review of Applied Science and Technology*, 6(1), 263-279. <https://doi.org/10.63125/cm1c2c20>
- [19]. Mst Shurovi, A. (2024). Machine Learning Models for Bank Reconciliation Anomaly Detection in High-Volume Financial Transactions. *American Journal of Data Science and Analytics*, 5(12), 314-354. <https://doi.org/10.63125/mmwb6467>
- [20]. Muhammad Zahidul, I. (2024). Assessment of Data Integrity Gaps and Preventive Strategies in QC Laboratories in Compliance with FDA and CGMP Expectations. *American Journal of Scholarly Research and Innovation*, 3(01), 243-291. <https://doi.org/10.63125/a3xjb608>
- [21]. Muhammad Zahidul, I. (2025). Development of a Dashboarding and Visualization Framework For Pharmaceutical Brand Performance Analytics Using Tableau and Python. *American Journal of Interdisciplinary Studies*, 6(3), 351-393. <https://doi.org/10.63125/e6dj308>
- [22]. Muhammad Zahidul, I., & Amena Begum, S. (2022). Applied Data-Analytics Approaches to Cardiovascular Drug-Utilization Patterns: An Empirical Study Using Real-World Pharmaceutical Sales Data. *American Journal of Data Science and Analytics*, 3(12), 89-129. <https://doi.org/10.63125/7xgryf71>
- [23]. Mukut Kanti, B. (2025). Lifecycle Carbon Reduction and Net Zero Pathway Modeling Using LEED, EDGE, and SBTI Frameworks in Developing Economy Industrial Buildings. *American Journal of Interdisciplinary Studies*, 6(3), 306-350. <https://doi.org/10.63125/tnneth75>
- [24]. Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection. *Decision Support Systems*, 50(3), 559-569. <https://doi.org/10.1016/j.dss.2010.08.006>
- [25]. Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206-215. <https://doi.org/10.1038/s42256-019-0048-x>
- [26]. Sharif Md Yousuf, B., Md Shahadat, H., Saleh Mohammad, M., Mohammad Shahadat Hossain, S., & Intiaz, P. (2025). The Green Hydrogen Blueprint: Optimizing America's Clean Energy Future. *American Journal of Data Science and Analytics*, 6(10), 43-69. <https://doi.org/10.63125/y9b1k470>
- [27]. Shima Ali, S., Musomi, K., & Sonya, G. (2024). Integrated Infectious Disease Preparedness: Multilingual Health Communication Strategies for Public Health Risk Mitigation, Disease Intervention and Prevention for Culturally Diverse Population in United States. *American Journal of Data Science and Analytics*, 5(10), 01-47. <https://doi.org/10.63125/ehk4mt02>
- [28]. Shima Ali, S., Sonya, G., & Mahbubul, I. (2025). Telemedicine And Remote Patient Monitoring for Preventing Risk and Chronic Disease Management in Underserved Communities. *American Journal of Health and Medical Sciences*, 6(04), 01-42. <https://doi.org/10.63125/y0e9dc38>

- [29]. Tonay, P., Md Maruf, I., & Al, A. (2024). Artificial Intelligence-Based Decision Support Systems and Managerial Performance. *American Journal of Advanced Technology and Engineering Solutions*, 4(04), 154-184. <https://doi.org/10.63125/wmz8dc67>
- [30]. Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., ... van Mulbregt, P. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3), 261-272. <https://doi.org/10.1038/s41592-019-0686-2>