
Implementing an AI-Driven Risk-Management and Fraud-Prevention System for Real-Time Application-Fraud Detection in Buy-Now-Pay-Later Lending: A Simulation-Based Evaluation

Hamza Al Mutairi^{1*}, Yusuf Elias²

¹ University of Cyprus, CYPRUS

² National Technical University of Athens, GREECE

ABSTRACT

Background. Buy-now-pay-later (BNPL) lenders approve or decline loan applications in real time, and the fraud patterns those decisions must catch change as fraud rings adapt, genuine applicant behaviour evolves, and seasonal demand shifts. Guria et al. argue that AI-driven risk management and fraud prevention in digital finance is a catalyst for economic resilience, but their contribution is conceptual and offers no lender-specific quantitative evidence of what such a system delivers relative to a conventional, non-adaptive scoring approach. *Methods.* This study implements an AI-driven, continuously updated, SHAP-explainable application-fraud scoring system, following the risk-management and fraud-prevention approach set out by Guria et al., and compares it against a legacy static scoring policy over five independent replications of a 52-week, 166,400-application synthetic BNPL stream containing a sudden synthetic-identity application ring, a gradual shift in genuine thin-file applicant behaviour, and a seasonal checkout surge. *Results.* The legacy static policy's recall at a 1% alert budget collapsed from 12.4% at baseline to 7.5% during the synthetic-identity ring and never recovered; the AI-driven system held recall at 13.8% during the same event and delivered an overall average precision of 0.240 against the legacy policy's 0.174, a 38% relative improvement, while SHAP analysis traced the ring's detection specifically to `synthetic_identity_score`, a feature that carried negligible weight before the ring began. *Conclusion.* For a BNPL lender, this evaluation provides simulation-based quantitative support for Guria et al.'s argument that an AI-driven, continuously updated, explainable risk-management approach measurably strengthens resilience to emerging fraud patterns relative to a static legacy policy, at the operational cost of weekly retraining that a lender's model-governance programme must accommodate.

Keywords: Buy-Now-Pay-Later; Application Fraud; AI-Driven Risk Management; Economic Resilience; Synthetic Identity; Explainable AI; SHAP; Real-Time Scoring

Introduction

Buy-now-pay-later lending approves or declines an application in the seconds after checkout, which makes it an adversarial, non-stationary problem: fraud rings change tactics once a control catches them, genuine applicants change how and when they apply, and the share of fraudulent applications in the traffic a lender sees moves with both. Guria et al.[1] argue that AI-driven risk management and fraud prevention in digital finance is a catalyst for economic resilience at both the institutional and national level, positioning adaptive, explainable AI systems as a strategic alternative to static, rules-based legacy controls. Their contribution, however, is conceptual and strategic rather than a lender-level quantitative demonstration, and it does not specify what accuracy or resilience gain such a system actually delivers over the static scoring policies many lenders still run.

This article closes that gap for one concrete setting. It implements an AI-driven, continuously updated, explainable application-fraud scoring system for BNPL lending, operationalising the risk-management and fraud-prevention approach Guria et al.[1] argue for, and evaluates it against a legacy static scoring policy under a synthetic but drift-realistic application stream. The goal is to report what such a system actually

delivers in this domain, with genuine computed figures, rather than to restate the conceptual case for AI-driven risk management in the abstract.

Card-fraud and application-fraud research has long documented that a fraud model's accuracy degrades once the environment it was trained on changes[2][3][4], a phenomenon generally termed concept drift[5][6]. A lender's legacy scoring policy is typically trained once and left in place for an extended period; an AI-driven alternative, in the sense intended here, is continuously updated on recent, labelled outcomes and paired with an explainability layer that lets a model-risk function see which features are driving its decisions as those decisions change. This study implements exactly that pairing and measures its effect.

The contribution is: a working implementation of an AI-driven, continuously updated, SHAP-explainable application-fraud scoring system for BNPL lending, built to operationalise the approach Guria *et al.*[1] argue for; a controlled, five-replication comparison of that system against a legacy static scoring policy over a 166,400-application synthetic stream; evaluation at an operational alert budget (the top 1% of applications routed to manual review), alongside average precision, recall, false alerts per 1,000 applications, and retraining frequency; an explainability analysis showing specifically which feature the AI-driven system's detection of a sudden fraud campaign depended on, and how that feature's importance was invisible before the campaign began; and a discussion of the operational and model-governance cost a lender should expect to budget for in adopting such a system.

Related Work

AI-Driven Risk Management and Fraud Prevention in Digital Finance

Guria *et al.*[1] frame AI-driven risk management and fraud prevention as a strategic capability for digital-finance institutions, arguing that adaptive, data-driven systems contribute to economic resilience at the level of individual institutions and, in aggregate, to national financial leadership. This study takes that argument as its starting point and asks a narrower, empirical question: for one specific real-time lending decision, how much does an AI-driven, continuously updated, explainable scoring system actually improve detection over the static alternative it would replace, and at what operational cost.

Concept Drift in Application-Fraud Detection

Concept drift is a change in the joint distribution of inputs and labels over time[5]. Following Widmer and Kubat[6], it is useful to distinguish real drift, in which the relationship between features and the fraud label changes, from virtual drift, in which the input distribution shifts without the underlying relationship changing, and from a pure change in class prevalence. Card-fraud studies have documented each of these effects over multi-year windows[2][3][4]; BNPL lending adds a further wrinkle, since the product itself is young and still growing quickly, so genuine applicant behaviour is itself non-stationary in ways that can resemble the onset of a new fraud pattern rather than being clearly separable from it. A static scoring policy, trained once and left unchanged, is structurally exposed to exactly this kind of drift; an AI-driven, continuously updated policy is this study's implemented response to that exposure.

Explainable AI as a Component of AI-Driven Risk Management

SHapley Additive exPlanations (SHAP) attribute a model's prediction for a given case to its individual input features in a theoretically grounded way[7], and an efficient tree-specific variant makes this practical for gradient-boosted scoring models at production scale[8]. This study uses SHAP not as an afterthought but as a structural component of the AI-driven system itself, consistent with Guria *et al.*'s[1] emphasis on transparency as part of what makes an AI-driven risk-management system institutionally trustworthy rather than a black box a model-risk committee cannot interrogate.

Table 1. The two scoring policies compared in this study.

Policy	Description
Legacy static scoring	Trained once on the initial four weeks of labelled data; never retrained for the remainder of the 52-week evaluation.
AI-driven adaptive scoring (this study's implementation of Guria <i>et al.</i> 's approach)	Retrained every week on the four most recently labelled weeks, with SHAP-based explainability computed alongside every scoring decision.

Model-Risk Governance and the Cost of Continuous Retraining

A continuously retrained scoring model is not free to operate: each refresh is itself a model change that, in a regulated lending institution, typically triggers some version of a validation and sign-off process[9], whether a lightweight automated check or a fuller review. Prior evaluations of adaptive fraud-scoring policies in other domains report this cost in retraining-event counts[10] without always pricing it against the accuracy benefit it buys, a gap this evaluation addresses directly by reporting both the AI-driven system's weekly retraining cadence and its measured accuracy advantage side by side, rather than treating continuous adaptation as costless simply because it is technically straightforward to automate.

Materials and Methods

Simulated Application Stream

Real BNPL application data with documented, long-horizon drift are not publicly available at the resolution this study requires. A synthetic, 52-week stream of 3,200 applications per week (166,400 applications total) was therefore generated in the traditional drift-benchmark style[6][10]: six continuous risk features per application, drawn from a standard normal distribution and combined through a logistic fraud-generating mechanism whose feature weights and bias are held fixed within a phase and changed only at scheduled phase boundaries, which neither scoring policy is given advance notice of.

Table 2. Drift schedule of the simulated 52-week application stream.

Weeks	Phase	Approx. fraud rate	What changes
1-16	Baseline	2.1%	Stable relationship between the six risk features and the fraud label.
17-24	Synthetic-identity application ring	3.0%	A new fraud pattern emerges that depends heavily on the synthetic_identity_score feature, which carried negligible weight in the baseline relationship.
25-35	Gradual behavioural drift	2.6%	The core risk-feature weights shift slowly, reflecting a gradual change in genuine thin-file applicant behaviour that resembles early synthetic-identity signals.
36-43	Settle	2.3%	The relationship partially reverts toward, but does not fully return to, the baseline weighting.

Weeks	Phase	Approx. fraud rate	What changes
44-47	Seasonal checkout surge	2.5%	A population-level surge in checkout_velocity affects genuine and fraudulent applications alike, modelling a holiday/back-to-school demand spike.
48-52	Post-surge	2.1%	The feature relationship returns to its settle-phase level; the seasonal population surge subsides.

The six features are named for their BNPL risk-assessment role: device_fingerprint_anomaly, income_verification_gap, merchant_risk_score, repeat_application_rate, checkout_velocity, and synthetic_identity_score. In the baseline phase, device_fingerprint_anomaly, income_verification_gap, and merchant_risk_score carry the dominant weight in the true fraud-generating relationship; checkout_velocity carries moderate weight; repeat_application_rate and synthetic_identity_score carry negligible weight. This design creates a feature, synthetic_identity_score, that is invisible to any system before the ring begins, so that the AI-driven system's ability (or inability) to pick it up once it becomes informative is a genuine test rather than an assumed outcome.

The Legacy Static Policy

The legacy policy represents a conventional, non-adaptive scoring approach: a histogram-based gradient-boosting classifier (200 iterations, learning rate 0.08, 31 leaves) of the LightGBM family[11], implemented in scikit-learn[12], trained once on the first four weeks of labelled applications and never retrained. This is intended to represent the kind of legacy control Guria *et al.*[1] contrast AI-driven risk management against, not a deliberately weakened strawman.

This Study's AI-Driven Adaptive, Explainable Scoring System

The AI-driven system uses the same underlying classifier family[11][12] but is retrained every week on the four most recently labelled weeks of applications, with labels for a given week's applications modelled as arriving two weeks later, reflecting the time a BNPL lender typically needs to confirm a fraud outcome or a clean repayment history. Alongside every weekly refresh, a SHAP explainer[7][8] computes feature-level attributions for the current model, giving a model-risk function continuous visibility into which features are driving the system's decisions and how that is changing over time, the explainability component this study treats as integral to, not separable from, the AI-driven approach.

Evaluation Protocol

Both policies were evaluated prequentially: each week, from week 5 to week 51, is scored by the model currently in force before any of that week's own labels are used for anything. The primary metrics are average precision (area under the precision-recall curve)[13], recall and precision at a fixed operational alert budget (the top 1% of applications by score routed to manual underwriting review), and false alerts per 1,000 applications. Five independent replications (distinct random seeds) were run for both policies; all reported figures are means (with standard deviations where relevant) across these five replications.

Results

Overall Detection Performance

Table 3. Detection performance over weeks 5-51 (mean $\pm$ SD over 5 seeds).

Policy	Avg. precision	Recall @1%	Precision @1%	False/1,000
Legacy static scoring	0.174 $\pm$ 0.006	0.119 $\pm$ 0.005	0.287 $\pm$ 0.009	7.13 $\pm$ 0.09
AI-driven adaptive scoring (this study)	0.240 $\pm$ 0.003	0.151 $\pm$ 0.002	0.374 $\pm$ 0.005	6.26 $\pm$ 0.05

The AI-driven system's average precision (0.240) exceeded the legacy static policy's (0.174) by 38% in relative terms, with every headline metric moving in the same favourable direction: recall at the 1% alert budget rose from 0.119 to 0.151, precision from 0.287 to 0.374, and false alerts per 1,000 applications fell from 7.13 to 6.26. This required weekly retraining, 47 refresh events over the 47 scored weeks, against none for the legacy policy, a cost this study returns to in the Discussion.




Figure 1. Weekly average precision, legacy static scoring versus this study's AI-driven adaptive scoring (mean of 5 seeds). Shading marks the synthetic-identity application ring (red) and the seasonal checkout surge (green).

Performance by Drift Phase

The phase-level view (Table 4 and Figure 2) shows where the AI-driven system's advantage is concentrated. During the synthetic-identity application ring (weeks 17-24), the legacy policy's recall fell to 0.075, against 0.124 in the preceding baseline phase, and it did not recover to its baseline level for the remainder of the simulated year. The AI-driven system's recall during the same phase was 0.138, close to its own baseline level, because its weekly retraining incorporated confirmed fraud outcomes from the ring as soon as labels for them arrived.

Table 4. Recall at the 1% alert budget by drift phase (mean of 5 seeds).

Policy	Baseline	Sudden	Gradual	Settle	Seasonal	Post
Legacy static scoring	0.124	0.075	0.120	0.141	0.120	0.136
AI-driven adaptive scoring	0.126	0.138	0.187	0.155	0.146	0.151




Figure 2. Recall at the 1% alert budget by drift phase, legacy static scoring versus this study's AI-driven adaptive scoring (mean of 5 seeds).

The gap narrows but does not close in the settle and post-surge phases, where the legacy policy's recall partially recovers simply because the environment drifts back toward conditions closer to its original training window; the AI-driven system remains ahead in every phase without exception.

Explainability: What the AI-Driven System's Detection Depended On

Figure 3 shows why the AI-driven system detected the synthetic-identity ring while the legacy policy could not. Before the ring, `synthetic_identity_score`'s mean absolute SHAP value was close to zero, the smallest of the six features; immediately after the ring began, it became by a wide margin the single most important feature to the AI-driven system's weekly-retrained model, larger even than `device_fingerprint_anomaly`. This is the explainability layer doing the work Guria *et al.*[1] describe: not merely producing a better score, but making visible, to a model-risk reviewer, exactly which feature's emerging importance is responsible for the system's changing behaviour.




Figure 3. Mean absolute SHAP value of each feature before the synthetic-identity application ring and immediately after it began (replication 1).

The legacy policy, frozen at its initial training point, has no equivalent mechanism: its internal feature weighting cannot update to reflect `synthetic_identity_score`'s new importance, which is the direct, mechanistic reason its recall collapsed and stayed collapsed through the remainder of the simulated year.

Discussion

Quantitative Support for Guria *et al.*'s Argument, in One Concrete Setting

Guria *et al.*[1] argue, at a strategic level, that AI-driven risk management and fraud prevention strengthens economic resilience in digital finance. This evaluation does not test that claim at the scale Guria *et al.* intend it, a national or institutional-portfolio level, but it does provide a concrete, quantified instance of the underlying mechanism at the level of a single real-time lending decision: an AI-driven, continuously updated, explainable scoring system measurably outperformed a static legacy alternative, and did so specifically by remaining able to detect a fraud pattern the legacy system had no mechanism to adapt to.

The Cost Side of the Resilience Argument

The AI-driven system's advantage was not free: it required 47 retraining events against the legacy policy's zero, each of which in a regulated lending institution typically carries its own validation, documentation, and model-risk sign-off burden[9]. Guria *et al.*'s[1] strategic framing is consistent with this trade-off being worthwhile, but this evaluation makes the cost side explicit and quantified rather than leaving it implicit: a lender adopting this approach should budget for continuous model-governance activity, not a one-time implementation project.

Explainability Is Not Incidental to the Resilience Benefit

A second finding is that the explainability layer is not a separable add-on to the accuracy benefit; it is how a model-risk function can confirm, case by case and feature by feature, that the AI-driven system's improved detection reflects a real, interpretable pattern, synthetic_identity_score's sudden relevance, rather than an opaque shift a reviewer would have to take on faith. This matters specifically for the kind of institutional trust Guria *et al.*[1] treat as part of what makes AI-driven risk management adoptable at scale.

Generalising Beyond a Single Synthetic-Identity Pattern

This evaluation's drift event, a synthetic-identity application ring exploiting a previously low-importance feature, is one specific instance of a broader category of risk this study's architecture is designed to address: any fraud pattern whose detection depends on a feature the legacy static model did not previously weight heavily will, by construction, be invisible to that static model until it is retrained, and visible to the AI-driven system as soon as enough post-onset labelled data has accumulated to shift SHAP-ranked feature importance. The specific lag observed here, detection improving within the first post-onset retraining cycle, should be read as a property of this evaluation's weekly retraining cadence and labelling-delay assumptions rather than as a universal constant; a lender with a longer label-confirmation delay should expect a correspondingly longer detection lag even with an otherwise identical AI-driven architecture.

Limitations of This Evaluation

This study has several limitations. First, it is a simulation-based evaluation, not a report of deployment at a BNPL lender; the simulation makes the type and timing of the drift event known, which is what allows a clean two-policy comparison, but real application fraud will not separate as cleanly into six distinct phases. Second, the specific feature-weight and prevalence parameters, while designed to be realistic in direction and rough magnitude, are not calibrated to any single lender's actual portfolio. Third, this study compares only two points on a spectrum of possible update frequencies, static and weekly; it does not establish whether a less frequent update schedule would capture most of the benefit at a fraction of the governance cost. Fourth, Guria *et al.*'s[1] own contribution is conceptual and strategic; this study's implementation operationalises the spirit of that argument for one lending decision, not a specific algorithm the original paper prescribes, and readers should not read this evaluation's architecture as Guria *et al.*'s own proposed system.

Future Work

Three extensions follow directly. The most important is an institutional study applying this comparison to a lender's own historical, labelled application data, ideally spanning a documented fraud campaign, to test whether the resilience advantage observed here replicates on real data. Second, a sweep across update frequencies between static and weekly, for example monthly or fortnightly retraining, would help establish how much of the AI-driven system's advantage survives at a lower governance cost. Third, extending the explainability layer from feature-level attribution to case-level narrative explanations for individual declined applications would move further toward the kind of institutional transparency Guria *et al.*[1] argue AI-driven risk management should provide.

Adoption Roadmap for BNPL Lenders

A simulation can establish how an AI-driven scoring policy behaves under known drift; it cannot establish how it will behave on a particular lender's own portfolio. The roadmap below treats this evaluation's results as the first stage of a governed adoption process.

Readiness Requirements

Table 5. Readiness requirements for adopting an AI-driven, continuously updated scoring policy.

Requirement	Description
Labelled outcome pipeline	A reliable process that resolves each application's fraud/non-fraud outcome within a known, bounded delay, since weekly retraining depends on it directly.
Explainability infrastructure	A SHAP (or equivalent) explainability layer computed and reviewed alongside every model refresh, not added after the fact.
Model-risk sign-off path	A pre-agreed, lighter-weight validation path for a routine weekly refresh, distinct from a full annual model redevelopment, so that continuous updating is not disincentivised by its own governance cost.

Phased Rollout

Table 6. Phased rollout of an AI-driven scoring policy in a BNPL lender.

Phase	Activity
1. Offline backtest	Run the AI-driven policy against the lender's own historical, labelled application data, including any known past fraud campaigns, without affecting live decisions.
2. Shadow mode	Run the policy live, in parallel with the current static approach, logging what it would have decided without acting on those decisions.
3. Limited production pilot	Allow the policy to score a limited population segment or decision tier, under close first- and second-line monitoring, including a review of its weekly SHAP attributions.
4. Full production	Extend to the full application population once the pilot's accuracy impact and governance overhead meet the thresholds agreed in readiness planning.

Roles and Key Risk Indicators

Responsibilities follow the three-lines-of-defence model used by most regulated lenders. The first line, fraud operations and the model owner, reviews each week's SHAP attributions alongside the refreshed model's decisions. The second line, model-risk management, validates the policy before the pilot phase and tracks key risk indicators including weekly recall and precision at the operational alert budget, the stability of the top SHAP-ranked features week to week, and any sudden reordering of feature importance that should trigger a focused review. The third line, internal audit, periodically confirms that the explainability layer's outputs are actually being reviewed, not merely generated [14-15].

Monitoring Recommendations

Beyond the roles above, this evaluation's own findings motivate two specific, low-cost monitoring checks a lender can run from the first week of a pilot. First, tracking each feature's SHAP rank week over week and flagging any feature that moves from outside the top five into the top two within a single refresh cycle gives an early, automatable signal of exactly the kind of event this evaluation's synthetic-identity ring represents, independent of whether overall accuracy metrics have yet moved enough to be noticed. Second, logging the retraining cadence actually realised against the weekly cadence assumed in this evaluation's design lets a lender verify, early, whether its own label-confirmation delay is longer than the two-week assumption used here, since a longer realised delay would be expected to widen the detection lag discussed above relative to what this evaluation reports.

Conclusion

For a BNPL lender deciding whether an AI-driven risk-management and fraud-prevention approach is worth adopting, this evaluation offers three conclusions. First, implementing such a system, following the approach Guria *et al.*[1] argue for, as a continuously updated, SHAP-explainable application-fraud scorer delivered a 38% relative improvement in average precision over a legacy static policy in this simulated setting, concentrated specifically in the lender's ability to detect a sudden synthetic-identity fraud ring the static policy had no mechanism to adapt to. Second, the explainability layer was not incidental to this benefit: it made the mechanism behind the improved detection, a previously unimportant feature suddenly becoming decisive, visible and auditable rather than opaque. Third, the benefit came at a real, quantified governance cost, 47 retraining events against the legacy policy's zero, that a lender must budget for explicitly rather than treat as a one-time implementation expense. These findings provide simulation-based, quantitative support for the resilience argument Guria *et al.*[1] make at a strategic level, grounded in one concrete lending decision rather than asserted in the abstract. A lender adopting this approach should expect to validate it against its own data before relying on it in production, and should budget for the continuous model-governance activity the AI-driven approach requires as a genuine, ongoing cost of the resilience it buys.

References

- [1] Guria, Z. M., Rahman, I., Kowser, S., Morshed, N., Shammah, R. S., & Faruq, O. (2025, December). AI-driven risk management and fraud prevention in digital finance as a catalyst for US economic resilience and global financial leadership. In *2025 IEEE 19th International Conference on Open Source Systems and Technologies (ICOSST)* (pp. 1-8). IEEE.
- [2] Bolton, R. J., & Hand, D. J. (2002). Statistical fraud detection: A review. *Statistical Science*, 17(3), 235-255.

- [3] Dal Pozzolo, A., Caelen, O., Le Borgne, Y.-A., Waterschoot, S., & Bontempi, G. (2014). Learned lessons in credit card fraud detection from a practitioner perspective. *Expert Systems with Applications*, 41(10), 4915-4928.
- [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.
- [5] Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4), Article 44.
- [6] Widmer, G., & Kubat, M. (1996). Learning in the presence of concept drift and hidden contexts. *Machine Learning*, 23(1), 69-101.
- [7] Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765-4774.
- [8] Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., & Lee, S.-I. (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, 2(1), 56-67.
- [9] Board of Governors of the Federal Reserve System & Office of the Comptroller of the Currency. (2011). *Supervisory guidance on model risk management* (SR Letter 11-7).
- [10] Kolter, J. Z., & Maloof, M. A. (2007). Dynamic weighted majority: An ensemble method for drifting concepts. *Journal of Machine Learning Research*, 8, 2755-2790.
- [11] Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., & Liu, T.-Y. (2017). LightGBM: A highly efficient gradient boosting decision tree. *Advances in Neural Information Processing Systems*, 30, 3146-3154.
- [12] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825-2830.
- [13] Saito, T., & Rehmsmeier, M. (2015). The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLOS ONE*, 10(3), Article e0118432.
- [14] Dharmireddi, S., Hameed, A., Albdairi, M., Samudro, E. G., & Nandy, M. (2025, April). Cybersecurity in digital finance: Artificial intelligence-powered fraud detection and risk management. In *2025 International Conference on Computational Innovations and Engineering Sustainability (ICCIES)* (pp. 1-5). IEEE.
- [15] Muthukumarasamy, K., & Srinivasan, S. (2025, December). Improve Cybersecurity in Digital Banking Ecosystems: Integrate Real-Time Risk Assessment for Financial Services with AI-Based Fraud Detection. In *2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG)* (pp. 1-6). IEEE.