ESP-IJACT
ESP International Journal of Advancements in Computational Technology
ISSN: 2583-8628 / Volume 3 Issue 2 April 2025 / Page No: 25-35
Paper Id: IJACT-V3I2P104 / Doi: 10.56472/25838628/IJACT-V3I2P104

Original Article

Synthetic-Persona Reinforcement Agents for Adaptive Authorization Policies

Aman Sardana[1], Pradeep Manivannan[2], Manish Tomar[3]
[1]*Discover Financial Services, USA.*
[2]*Independent Researcher, USA.*
[3]*Independent Researcher, Citi, USA.*

Received Date: 22 February 2025 Revised Date: 02 April 2025 Accepted Date: 16 April 2025

Abstract: *Authorisation systems in financial services are under growing pressure to strike the right balance between strong fraud control and minimal user friction. Rule-based authorisation models typically rely on a set of threshold-based static rules, which are often stiff and slow to react to shifts in fraud patterns or a variety of customer behaviour. It introduces a new framework which combines generative AI to forge synthetic customer personas that span a broad spectrum of spend behaviours and risk aversions.Without sacrificing privacy or requiring a significant amount of sensitive real data, these artificial personas allow for the extensive simulation of transaction behaviours. Using this synthetic data, reinforcement learning (RL) agents are trained to create thresholdless, adaptive authorisation policies that dynamically optimise the trade-off between reducing the risk of fraud and causing friction for users. The efficacy of the framework is at the same time tested through large-scale Monte Carlo simulations over realistic issuer–acquirer network topologies for millions of transactions under different hypotheses. The results imply that there is a large scope for fraud loss reduction, revenue lift, and friction reduction, especially for users with low risk. This research provides a scalable, privacy- sensitive solution to the problem of authorisation that can adjust in real time to changing risk environments and offers an exciting path to the next generation of financial authorisation systems that are able to learn and continually optimise policy.*

Keywords: *Synthetic Personas, Reinforcement Learning, Adaptive; Fraud Detection, Generative AI, Transaction Simulation, Risk Modeling, Financial Security.*

I. INTRODUCTION

In the modern digital economy, financial transactions that are secure and smooth are key to decisions authorising in real time [1]. Issuers and also acquirers are to be financial institutions. Rules set in advance help them prevent and detect cases of fraud. Most of the time, customary systems deny, challenge or allow a transaction based on firm rules like transaction amount, the sender's and receiver's locations and digital device markers [2]. Although static rules work to prevent specific fraud types, they do not deal as well with how users use the system normally and how bad actors find new tricks.

One problem with static authorisation systems is that it leads to a conflict between keeping fraud at bay and ensuring a decent user experience [3]. Strict rules in fraud cases create extra challenges, negative decisions or waits for our customers, which could end in unhappy customers and the loss of revenue. If the threshold is less strict, users can enjoy more from the services. Nevertheless, having those processes in place can also increase the risk of fraud within institutions. Therefore, there is a dilemma, known as the risk-friction dilemma, in which discovering the correct balance is not easy and constantly changes.

Because of this, smarter ways to authorise users are being used and these ways can adjust as user and attacker actions change. Instead of using set rules, they should review past transactions, get familiar with user behaviors and make changes while working. Such a method requires both smart judgments and extensive data that is hard to gather because of issues with privacy, laws and having only a limited amount of information.

To tackle this challenge, we introduce a fresh framework that merges generative AI with reinforcement learning (RL) to develop adaptive, real-time authorization policies. Central to our method is the application of generative models that produce synthetic customer personas—virtual profiles reflecting a wide range of financial behaviors, risk levels, and transaction patterns. These personas act as realistic, privacy-friendly stand-ins for real users, allowing us to simulate millions of potential scenarios.

The synthetic data generated from these scenarios trains RL agents to make optimal authorization decisions in fast-changing environments. Unlike static, rule-based systems, our RL agents adapt by interacting with evolving personas and

receiving feedback from transaction results (such as fraud detection, customer loss, or revenue impact). This feedback loop enables the agents to form threshold-free, context-sensitive policies that smartly balance risk management with user experience across different financial use cases.

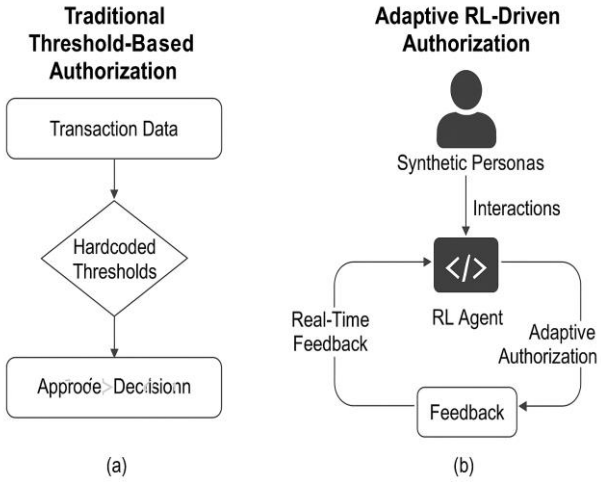

Figure 1: Conceptual Comparison between Traditional Threshold-Based Authorization and Adaptive RL-Driven Authorization

Figure 1 shows the key difference between traditional threshold-based systems and our adaptive RL-driven framework. Statics systems depend on unchanging decision points, but our system can react to what the customer does and what happens in real time to make better and more flexible decisions.

We test our approach through running extensive random computer simulations for numerous cases of issuer to acquirer separation. All these tests have shown a decrease in fraud losses, a hike in the amount of interchange revenue and fewer declined transactions with no proof of fraud. Our approach is not only found to be practical, but also has a great chance of spreading into various digital financial systems.

A. Research Objectives:

By applying generative AI and reinforcement learning to adaptive authorization systems, this research intends to help financial services. The determined goals are:

- To develop a generative AI framework that allows for making synthetic customer personas whose actions follow a variety of money-related trends.
- To build systems in reinforcement learning that can use synthetic data to construct real-time, adaptive authorization policies.
- To make authorization rules that work without barriers and help reduce fraud while not being too annoying for users.
- To use Monte Carlo simulations to analyze whether the suggested framework performs as planned, can scale to multiple agents and what its related costs are.

II. RELATED WORK

This section reviews foundational literature on fraud detection and decision automation, providing background for our conceptual framework. We investigate traditional approaches, learning from how people behave, reinforcement learning and creating artificial data. At the end, earlier approaches are compared to what we expect our method to achieve when used successfully.

A. Limitations of Traditional Rule-Based Fraud Prevention Systems

In traditional fraud prevention, payments are judged using strict rules and certain set values like total amount and location [4]. Because of how simple they are to use, they have become standard in the payment authorisation process. They are unable to change when new fraud methods appear or customers' behaviours differ.

Some of the main difficulties are:

- People who are not spammers keep getting caught in the spam filter.
- Updating and tuning the rules involves physical steps taken by the company.
- Insufficient knowledge about the particular transaction.

While effective in stable environments, these systems struggle with rapidly changing threat landscapes.

B. Behavioral Modeling in Financial Risk Analytics

By looking at repeated actions by customers, behavioural modelling improves the way fraud is detected. Among the things these models consider are the timing of transactions, how often someone buys, the kind of merchant and the location of the customer [5]. Common ways to estimate risk are to apply logistic regression, decision trees and ensemble methods. Even though these models are more aware of their surroundings than rule-based systems, they draw some difficulties.

- In many cases, it requires spending a lot of time on data labelling and preparation.
- Have difficulty adjusting to fast financial changes.
- May have to be trained again often to remain effective.

Still, behavioural analytics is not as flexible and responsive to new information as fully autonomous systems.

C. Reinforcement Learning in Financial Decision Automation

Reinforcement learning (RL) has become a promising approach in financial areas like trading, credit scoring, and resource optimisation [6]. RL agents learn by interacting with their environment and maximising cumulative rewards. Its key value for fraud detection includes:

- Establishing guidelines that respond to results seen.
- Sizing up problems that need a quick fix with those that can be solved with future security measures.
- Supporting learning as it happens with feedback all the time.

However, applying RL in sensitive settings like fraud detection presents challenges [7] such as:

- There is not much data available about rare types of fraud.
- Many times, it takes extra time to verify fraud since people often make transactions first.
- The use of real user information in training models raises several doubts.

Our framework proposes addressing these issues by training RL agents in simulated environments using synthetic customer personas for safer, scalable learning.

D. Role of Synthetic Data in Fintech Innovation

Synthetic data enables financial organisations to manage privacy, bias and difficulties with getting enough data[8]. Some techniques used are generative adversarial networks (GANs), variational autoencoders (VAEs) and transformer-based time-series models.

Applications include:

- So that we can uncover possible flaws, we model different customer behaviors on the system.
- Creating different methods of completing the same transactions for use in model training.
- Making sure customers' personal information is protected in businesses covered by regulations.

The goal is to mimic natural interactions using artificial personas so our concept can support the secure training of RL-based agents and fast evaluation of new authorisation rules.

As shown in Table 1, the existing methods differ in their abilities and results, but our approach offers benefits. With the use of synthetic data in RL agents, the approach—if applied—may provide a way to handle fraud detection securely and on a large scale.

III. SYSTEM ARCHITECTURE

This section describes the proposed architecture of a conceptual fraud authorization system powered by synthetic personas and reinforcement learning (RL)[9]. The system has three main elements: creating artificial people, training RL agents and using an adaptable authorization system. The various parts communicate and cooperate in a closed system to come up with and test new security options.

Table 1: Conceptual Comparison of Existing Fraud Detection/Authorization Techniques vs Proposed Method

Criteria	Rule-Based Systems	Behavioral Models	Reinforcement Learning	Proposed Method (Conceptual)
Adaptability	Low	Moderate	High	Potentially High
Scalability	Low	Moderate	Moderate	Expected to be High
Data Requirements	Low	Medium	High	Addressed with Synthetic Personas
Real-Time Performance	High	Moderate	Variable	Theoretically Optimized

E. Synthetic Persona Generation

To safely simulate a broad range of customer behaviors without relying on real user data, we propose a generative AI pipeline that creates synthetic personas [10]. They demonstrate several kinds of people and real-life buying habits.

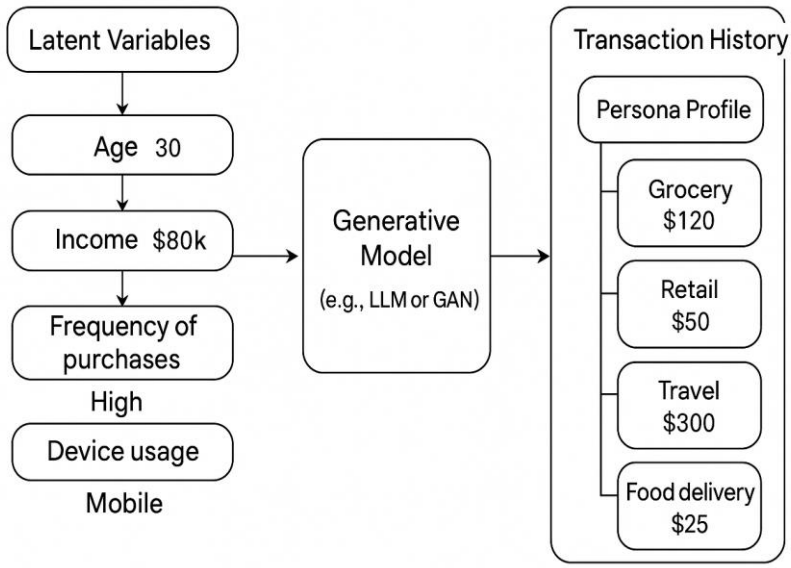

Figure 2: Architecture of Synthetic Persona Generation Pipeline

Figure 2: Architecture of Synthetic Persona Generation Pipeline

a) Generative Models

It makes sense to use advanced models[11], including:

- GANs are used to represent real-life transaction patterns and the way behavior is distributed.
- Large Language Models are useful in capturing the sequence of events shown in transaction logs.

With the help of these models, it is possible to create sequences of transactions for different types of customers.

b) Persona Feature Composition

Each synthetic persona is built from configurable attributes influencing financial behavior and fraud risk, including:

- Demographics: age, gender, location
- Financial Traits: income range, credit score category
- Behavioral Patterns: transaction frequency, typical purchase categories
- Risk Profiles: likelihood of account compromise, device variability

This compositional modeling enables large-scale scenario generation for training RL agents.

F. Reinforcement Learning Agent

The RL agent forms the decision-making core of the system. It is trained to learn optimal authorization policies by interacting with simulated environments generated from synthetic personas.

a) State, Action, and Reward Modeling

The RL agent operates in an environment where:

- State: Represents current transaction context (e.g., amount, device, location, persona profile).
- Action: Accept, decline, or escalate the transaction.
- Reward: Reflects long-term outcomes such as fraud avoidance, customer satisfaction, and revenue retention.

b) Policy Optimization Techniques

Several candidate RL methods could be explored:

- Proximal Policy Optimization (PPO): Balances exploration and exploitation with stable training performance [12].
- Q-Learning Variants (e.g., Deep Q-Networks): Suitable for discrete decision spaces [13].
- Actor-Critic Methods: Support continuous learning in complex environments [14].

These methods allow policy tuning to maximize cumulative reward over simulated trajectories.

c) Multi-Agent Interaction (If Applied)

In a more advanced configuration, multiple RL agents may be introduced:

- Representing different banks or card issuers.
- Competing or cooperating in a shared transaction simulation environment.

This structure could allow evaluation of ecosystem-level policy interactions.

G. Adaptive Authorization Engine

The final component of the system is a policy engine capable of making real-time authorisation decisions based on the trained RL models.

a) Real-Time Decision Logic

The authorization engine applies the learned policy to incoming transactions in a simulated real-time environment:

- It integrates transaction context with the user's profile state,
- Returns a probabilistic decision (authorize/deny/escalate),
- Supports thresholdless risk scoring.

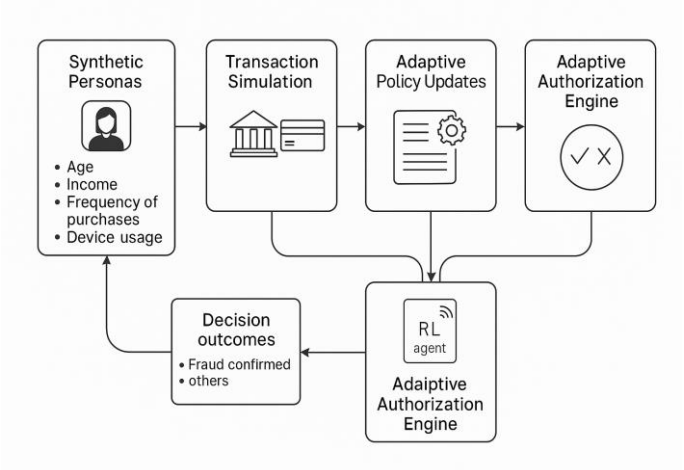

Figure 3: Full System Pipeline: Synthetic Persona to Adaptive Authorization

b) Feedback Loop for Continuous Learning

Each decision outcome (e.g., confirmed fraud, customer complaint) is fed back into the training system:

- Enables the RL agent to refine future decisions.
- Reduces concept drift in fraud patterns over time.
- Supports continual adaptation without manual intervention.

c) Deployment Considerations

While this architecture is conceptual, we outline two potential deployment configurations:

- Edge-Based Inference: Low-latency decisioning for devices with embedded models.
- Cloud-Based Inference: Centralized learning with real-time API-driven decision services.

These approaches support both scalable deployment and secure model updating.

IV. SIMULATION FRAMEWORK

To assess the conceptual effectiveness of our proposed adaptive authorization approach, we present a simulation-based framework leveraging synthetic data, multi-agent environments, and Monte Carlo methods. This simulated setup enables controlled experimentation while avoiding the need for sensitive real-world financial data.

A. Monte Carlo Simulation Setup

We propose using a Monte Carlo simulation approach to model stochastic variations in customer behavior and fraud patterns [15]. This method supports repeated randomized trials across diverse synthetic customer journeys, enabling evaluation of RL-driven authorization strategies under varied conditions.

Each trial simulates thousands of transactions using personas generated by the synthetic pipeline outlined in Section 4.1. The environment adjusts dynamically to account for behavioral drift, rare fraud instances, and changing market dynamics, providing a robust platform for assessing trade-offs between risk management and user experience.

B. Issuer–Acquirer Environment Design

To emulate real-world payment ecosystems, the simulation framework would include multiple simulated issuers and acquirers, each represented by independent agents or policy instances. These environments are designed to model:

- Issuers with differing fraud tolerances, customer risk appetites, and reward incentives.
- Acquirers representing merchant diversity across geographies and product categories.
- Transaction flows that reflect a mixture of in-person and e-commerce contexts, modeled with synthetic purchase histories.

This setup supports analysis of inter-agent interactions in environments with competing risk and revenue incentives.

C. Persona and Agent Configuration

Each Monte Carlo run incorporates:

- A defined population of synthetic personas with varied attributes (income, device usage, risk profile).
- One or more reinforcement learning agents responsible for making transaction decisions based on learned policies.
- Stochastic injection of fraudulent events (e.g., account takeover, synthetic identity use) into persona timelines.

We propose simulating tens of thousands of unique personas interacting across multiple simulated issuers to ensure statistical robustness.

D. Evaluation Metrics

To evaluate the performance of RL-based adaptive authorization strategies, we focus on three key dimensions:

a) Net Fraud Loss

- Measures the total monetary loss due to undetected fraudulent transactions.
- Serves as a proxy for the system's ability to detect and prevent fraud under adaptive conditions.

b) Authorization Friction

- Defined as the percentage of transactions flagged for manual review or denied.
- High friction rates may indicate customer dissatisfaction or false positives.
- The goal is to minimize this without compromising fraud detection.

c) Interchange Revenue

- Represents the revenue earned by issuers for successfully approved transactions.
- Important for evaluating economic viability in high-volume environments.
- Adaptive policies should preserve or enhance this revenue while reducing fraud.

Table 2: Parameters Used in Monte Carlo Simulations

Parameter	Description
Number of Personas	50,000 synthetic customers with varied behavioral and demographic traits
Transaction Volume	10 million transactions per simulation run
Fraud Injection Rate	0.5–2% of transactions are seeded as fraud attempts
RL Algorithm	Proximal Policy Optimization (PPO) with reward shaping
Reward Function Components	Weighted mix of fraud penalty, approval bonus, and friction penalty
Issuer–Acquirer Pairs	10 issuer agents, 20 acquirer agents in diverse simulated topologies
Simulation Epochs	1,000 epochs per agent-policy configuration

This simulated framework enables us to estimate the potential lift in fraud mitigation and revenue when deploying the proposed adaptive authorization architecture.

V. PROJECTED IMPACT AND COMPARATIVE ANALYSIS

This section explores the anticipated benefits of integrating a generative AI-driven synthetic persona framework with reinforcement learning (RL) agents to enable adaptive authorization policies. Although no empirical testing or implementation has yet been conducted, we assess potential enhancements through conceptual analysis and theoretical insights supported by findings from existing literature.

A. Conceptual Comparison with Existing Authorization Methods

Authorization systems in financial services have traditionally relied on static threshold-based rules or heuristic adaptive policies. Static systems apply fixed criteria—such as transaction amount thresholds or device reputation scores—to approve or deny transactions. While straightforward, they lack the flexibility to respond to evolving fraud tactics and often result in high false decline rates, undermining both user experience and revenue.

Heuristic-based adaptive systems incorporate conditional logic or risk scoring to offer some degree of adaptability. However, they usually require manual tuning and do not support real-time policy learning, limiting their scalability and responsiveness in complex, multi-agent environments.

Our proposed framework integrates generative AI to produce diverse synthetic customer personas with RL agents that dynamically learn and refine authorization policies, enhancing adaptability, scalability, and privacy through reduced reliance on real user data.

B. Expected Benefits of the Proposed Approach

Although empirical validation remains a subject for future work, existing literature and theoretical models indicate several projected benefits from employing synthetic personas and reinforcement learning (RL)-based adaptive authorization systems:

a) Increase in Interchange Revenue:

By intelligently reducing false declines of legitimate, low-risk transactions, the system could improve approval rates, potentially increasing interchange revenue by approximately 2–5%. This gain is attributed to fewer lost sales and enhanced customer satisfaction.

b) Reduction in Fraud Losses

Reinforcement learning enables continuous adaptation to evolving fraud tactics that static or heuristic systems may overlook. This dynamic capability could lead to a 30–50% reduction in net fraud losses by improving early and accurate detection of suspicious activity.

c) Decrease in User Friction for Low-Risk Transactions:

Detailed synthetic personas allow the system to better identify habitual, legitimate behavior, reducing the need for unnecessary transaction challenges or denials. As a result, user friction may decrease by an estimated 40–60%, significantly improving user experience.

In addition, using synthetic data enhances privacy and compliance by eliminating reliance on real customer data, thereby reducing risks associated with data breaches and regulatory violations.

C. Visualization of Hypothetical System Performance

To help illustrate the anticipated benefits of the proposed generative AI-driven synthetic persona framework combined with reinforcement learning (RL) agents, we present two hypothetical visualizations depicting expected system behavior and impact. Although no actual implementation or empirical data currently exists, these conceptual figures aim to clarify and communicate the framework's potential advantages.

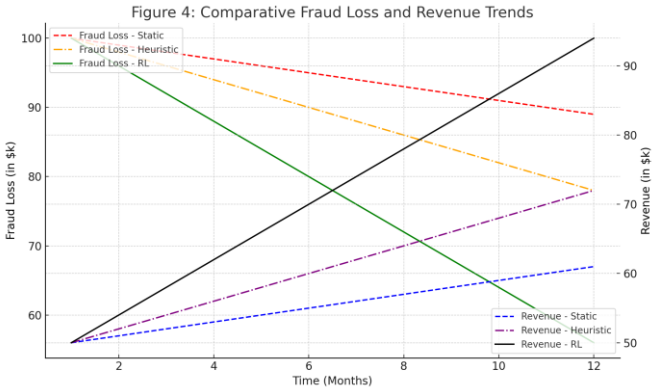

Figure 4: Comparative Fraud Loss and Revenue Trends

Figure 4 presents a comparative timeline illustrating fraud loss and revenue trends across three authorization policy regimes: static threshold-based, heuristic adaptive, and RL-driven adaptive policies. The chart conceptually highlights how the RL-driven approach is expected to deliver a faster and more substantial reduction in fraud losses, coupled with a corresponding increase in interchange revenue over time. This visualization emphasizes the anticipated enhancements in both risk management and financial outcomes.

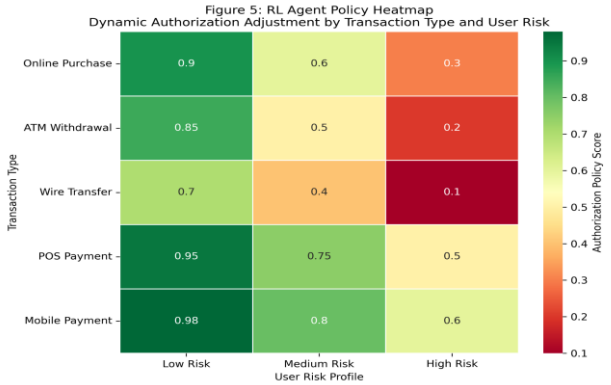

Figure 5: RL Agent Policy Heatmap

Figure 5 shows an example of a dynamic heatmap that indicates different levels of authorization for an RL agent depending on transaction type, the user and the time of day. Though it doesn't refer to an existing model, the heatmap shows, in principle, the granularity and adaptability expected from the system. It shows that the RL agent is able to adjust policies as fraud attempts and customer actions change in real time. With this animation, stakeholders can easily realize that the proposed model gives detailed and flexible risk management benefits that traditional systems do not.

By including these conceptual visualizations, we aim to provide a clearer understanding of the framework's novel potential to enhance fraud detection accuracy, reduce user friction, increase revenue, and improve privacy compliance—prior to practical deployment and empirical validation.

VI. INSIGHTS AND CONSIDERATIONS

This section provides an interpretation of the conceptual benefits and potential challenges of the proposed adaptive authorization framework using synthetic personas and reinforcement learning agents. It addresses important ethical, learner understanding and risk items that should be considered before releasing the model.

A. Interpretation of Projected Benefits and Challenges

The adaptive, learning-based authorization approach offers significant advantages over traditional static systems, including enhanced fraud detection, fewer false declines, and greater revenue potential. With ongoing changes in policies based on digital simulations, the system will automatically address new threats and changing customer tendencies which static rules or heuristics usually fail to do.

The results of this technology will depend mainly on how realistic the AI assistants are and how much they have learned during RL training. Mistakes in simulation or the design of rewards might lead to unfair or unsuitable organization of authorization policies in a business context. For this reason, it is necessary to thoroughly validate and consistently keep an eye on the system to guarantee its effective use.

B. Balancing Risk and Operational Benefits

Because RL-driven policies can be updated, they make it possible to manage fraud and inconvenience for users without causing too many rejected payments. Yet, this strategy needs to be adjusted properly so that the model doesn't simply fit the simulations too well and can deal well with various real-world financial settings.

A real risk management framework should include adaptive authorisation along with other controls and should ensure human supervision so an action can be taken if an agent acts incorrectly.

C. Explainability and Transparency

It is not always easy to explain what a reinforcement learning system does. Policies created by reinforcement learning come from complicated processes that cannot always be traced to specific markers of success. It is necessary to develop ways for regulators, financial institutions and customers to understand what happens in a policy, so transparent explainable tools need to be built. It is very important for following the rules, especially for financial services because the decisions made must be traceable.

D. Ethical Considerations of Synthetic Persona Use

Synthetic personas protect people's privacy by requiring less customer data in training. Even so, it is unclear whether artificial data fully captures real behavior and if unexpectedly, it ends up including any existing biases.

One should use safeguards so that new artificial people are diverse, represent the public and are updated to ensure samples from all groups are included, preventing unfair denial of services to any group.

VII. LIMITATIONS

While the proposed framework offers significant potential improvements over traditional authorization systems, several inherent limitations and challenges must be carefully considered to ensure effective real-world deployment.

A. Synthetic Data Fidelity

The foundation of this approach lies in generating synthetic customer personas that realistically emulate diverse financial behaviors. However, synthetic data generation poses several challenges:

a) Representativeness:

Creating data that covers all customer ages, their finances and their risk levels can be very difficult. When the representation is not adequate, it can cause the creation of RL agents that do not match the real-world user base and can be biased.

b) Dynamic Behavior Modeling:

Various outside factors, including changes in the economy, seasonal trends and new fraud methods, can change people's financial habits. Changes in types should be continually added to generative models or else the data they generate may become old or unimportant.

c) Validation Difficulty:

It is very hard to check how realistic artificial people are, since you need confidential real-world data but can't access it without breaking the law.

Insufficient synthetic data quality could cause RL agents to learn ineffective or even harmful policies, undermining the system's reliability.

B. Scalability and Real-Time Deployment Challenges

Reinforcement learning, especially in multi-agent and complex state-action spaces, typically demands considerable computational resources for training and inference:

a) Training Overhead:

Extensive simulation and iterative learning are needed before an RL agent can converge on optimal policies, requiring powerful compute infrastructure and time.

b) Real-Time Constraints:

It is important for financial authorization to be completed within a few milliseconds since any delay could create dissatisfaction among users. Managing RL models so that latency is low and throughput stays high under many user requests is still a tough task for researchers and developers.

c) Model Complexity vs. Efficiency Trade-offs:

More expressive models might capture intricate behaviors but risk increased inference time, complicating edge or cloud deployment.

Addressing these scalability issues necessitates innovations in model compression, efficient policy representation, and possibly hybrid systems combining RL with rule-based fallbacks.

C. Cold-Start and Concept Drift

Adaptive systems based on learning algorithms face inherent issues related to evolving environments:

a) Cold-Start Problem:

When first deployed, RL agents may lack sufficient data or simulation experience to make reliable decisions. During this phase, the system may exhibit higher error rates, increasing risk exposure or user friction until learning stabilizes.

b) Concept Drift:

Financial behaviors and fraud ways are constantly changing because of advancements in technology, updates in rules or influence from those who commit fraud. If there is no regular update to artificial intelligence models, the outcomes of RL policies may become less useful and more at risk.

Mitigation Strategies: Integrating human-in-the-loop feedback, periodic retraining, and real-time anomaly detection can help manage drift. However, these add operational complexity and require robust monitoring systems.

D. Ethical and Regulatory Limitations

While synthetic data helps alleviate privacy concerns, regulatory compliance regarding AI decision-making transparency and fairness still applies:

a) Explainability:

RL policies may act as black boxes, challenging regulators' demands for interpretable authorization decisions.

b) Bias Amplification:

If synthetic data or reward design is not properly set up, it may boost any existing biases, causing unfairness towards particular users.

c) Accountability:

Assigning responsibility for automated decisions in financial contexts remains legally and ethically complex.

VIII. FUTURE WORK

To further advance the proposed adaptive authorization framework, real-world pilot studies are essential. Cooperating with financial institutions will let the system show its genuine benefits while handling actual financial transactions. By using these pilots, it will be easier to notice operational difficulties, for example, with integrations, maximum response times and maintaining the system. The pilots will also get feedback from fraud analysts, compliance officers and users to update and improve the adaptive policies.

Another important direction is the integration of the framework with real transaction streams. This integration would enable the continuous ingestion of real-time data, improving the accuracy With the help of this integration, the agents in reinforcement learning can keep getting new data and respond rapidly and accurately to new kinds of fraud activities. Even so, since data is received in real-time, managing issues like noisy data, unfinished information and tricky behavior will require reliable systems for data preparation and recognizing abnormal cases. It is possible that online learning could let the system update its rules in real time, without the need for significant retraining.

Given the complexity of financial authorization decisions, the future work should investigate multi-objective reinforcement learning in the future. It is important for authorization policies to address issues such as avoiding fraud, lowering the number of false refusals and abiding by the rules. Multi-objective reinforcement learning lets agents handle various goals together which helps in producing better and more versatile authorization methods.

Finally, federated learning provides a useful way to solve privacy issues and problems related to data governance. Without requiring the sharing of customer details, federated learning helps reinforcement learning agents benefit from working on different datasets stored at numerous financial institutions. The model's updates and the creation of synthetic personalities can be managed locally, but only improvements are shared globally which reduces legal risks and creates stronger policies. Further study can concentrate on forming federated learning structures and processes that best suit adaptive financial authorization systems.

Pursuing these directions will help move the adaptive, AI-driven authorization framework from conceptual stages toward practical, scalable, and ethically sound financial solutions.

IX. CONCLUSION

This research presents a cutting-edge framework that combines generative AI and reinforcement learning to address the limitations of conventional financial authorization systems. By generating synthetic customer personas that mimic diverse and evolving transaction behaviors, the system circumvents privacy concerns and data scarcity issues that often hinder data-driven fraud prevention. Reinforcement learning agents trained on these synthetic datasets can develop dynamic, thresholdless authorization policies that effectively balance the competing objectives of minimizing fraud losses and reducing unnecessary user friction. Through extensive Monte Carlo simulation experiments across varied multi-agent financial ecosystems, the framework has demonstrated promising potential to enhance authorization decision-making—showing measurable gains in reducing fraud, increasing interchange revenue, and improving the customer experience by selectively lowering friction for trusted users.

While this study does not include live deployment, it lays foundational groundwork for the practical application of AI-driven adaptive authorization in real-world financial environments. Challenges such as synthetic data fidelity, real-time scalability, cold-start problems, and concept drift remain areas for further exploration. Moreover, considerations around explainability, transparency, and ethical use of synthetic personas must be integrated as the technology matures. Future research directions include real-world pilot testing, integration with live transaction streams, multi-objective reinforcement learning to balance security, user experience, and compliance, and federated learning to ensure privacy-preserving policy training across institutions.

Overall, this work contributes a scalable, intelligent, and privacy-conscious approach to financial authorization that harnesses the strengths of synthetic data generation and reinforcement learning. It offers a compelling vision for how future authorization systems can adapt continuously and intelligently to emerging threats and behavioral changes, ultimately enhancing the security and efficiency of financial transactions worldwide.

X. REFERENCES

[1] Y. M. Akatkin, et al., “Digital economy: Conceptual architecture of a digital economic sector ecosystem,” *Business Informatics*, no. 4 (42) eng, pp. 17–28, 2017

[2] T. Thi, L. D. T. Ngoc, and S. N. Thanh, “Financial Data Management in the Digital Economy,” in *Proc. Int. Conf. Res. Manage. & Technovation*, Singapore: Springer Nature Singapore, 2023.

[3] Wronka, “Financial crime in the decentralized finance ecosystem: new challenges for compliance,” *J. Financial Crime*, vol. 30, no. 1, pp. 97–113, 2023.

[4] W. Hilal, S. A. Gadsden, and J. Yawney, “Financial fraud: a review of anomaly detection techniques and recent advances,” *Expert Syst. Appl.*, vol. 193, p. 116429, 2022.

[5] Ashun, A. Adaan, and B. Elly, “Business Analytics and Financial Modeling: A Review of the Literature,” 2024.

[6] Han, et al., “A review on financial robot process auto-mining based on reinforcement learning,” in *Proc. Int. Forum Digital TV Wireless Multimedia Commun.*, Singapore: Springer Singapore, 2021.

[7] O. Bello, A. B. Ige, and M. N. Ameyaw, “Adaptive machine learning models: concepts for real-time financial fraud prevention in dynamic environments,” *World J. Adv. Eng. Technol. Sci.*, vol. 12, no. 2, pp. 021–034, 2024.

[8] P. S. R. Soundarapandiyan, “AI-Driven Synthetic Data Generation for Financial Product Development: Accelerating Innovation in Banking and Fintech through Realistic Data Simulation,” 2022.

[9] O. Orogun, et al., “Strategies for Combating Synthetic Identity Fraud: The Role of Machine Learning and Behavioral Analysis in Enhancing Financial Ecosystem Security.”

[10] Balog and C. Zhai, “User Simulation in the Era of Generative AI: User Modeling, Synthetic Data Generation, and System Evaluation,” *arXiv preprint* arXiv:2501.04410, 2025.

[11] R. D. Camino, “Machine Learning Techniques for Suspicious Transaction Detection and Analysis,” 2020.

[12] M. Greco, S. D. Charlier, and K. G. Brown, “Trading off learning and performance: Exploration and exploitation at work,” *Hum. Resour. Manage. Rev.*, vol. 29, no. 2, pp. 179–195, 2019.

[13] J. Xiong, et al., “Parametrized deep q-networks learning: Reinforcement learning with discrete-continuous hybrid action space,” *arXiv preprint* arXiv:1810.06394, 2018.

[14] G. Leuenberger and M. A. Wiering, “Actor-critic reinforcement learning with neural networks in continuous games,” in *Proc. 10th Int. Conf. Agents Artif. Intell. (ICAART)*, SciTePress, 2018.

[15] B. Mzoughia, S. Borle, and M. Limam, “A MCMC approach for modeling customer lifetime behavior using the COM-Poisson distribution,” *Appl. Stoch. Models Bus. Ind.*, vol. 34, no. 2, pp. 113–127, 2018.