



# Dynamic Fraud Detection in UPI Transactions

**L.SaiSampanPatrudu, D.Noorisha, E.Sidharatha, V.SaiRagaSudha**, students of  
Department of CSE-(AI&ML), Raghu Institute of Technology, Dakamarri(v),  
Bheemunipatnam, Visakhapatnam Dist, Pin Code: 531162

**Dr K. V. Satyanarayana, M. Tech (Ph.D), Professor, Head of the Department of CSE-  
(AI&ML), Raghu Institute of Technology, Dakamarri(v), Bheemunipatnam, Visakhapatnam  
Dist, Pin Code: 531162**

## ABSTRACT:

The paper titled "Dynamic Fraud Detection in UPI Transactions" proposes a novel approach to tackle the growing instances of fraudulent activities in digital payment systems, with a specific focus on Unified Payments Interface (UPI) transactions. The method employs Deep Reinforcement Learning (DRL) to adapt to the evolving fraudulent patterns on a dynamic basis, without the need for labelled data. This technique employs unsupervised learning methods to identify intricate fraudulent behaviours by iteratively engaging with transactional data. The research process is initiated by collecting and preprocessing data, followed by exploratory data analysis to identify patterns within the data. Subsequently, various machine learning algorithms, including logistic regression and gradient boosting, are applied to the preprocessed data for fraud detection. Additionally, the study innovates by developing a custom reinforcement learning environment, named FraudDetectionEnv, built on the OpenAI Gym framework. This environment trains reinforcement learning agents to detect fraud effectively. The integration of machine learning models into this reinforcement learning environment empowers agents to make informed decisions and refine their fraud detection strategies based on the accuracy of their actions. The proposed models' effectiveness is evaluated using real-world UPI transaction data, with key metrics such as accuracy, precision, and AUROC. The results demonstrate the method's superiority over traditional fraud detection approaches. This study underscores the potential of DRL in enhancing the security of digital payment ecosystems, suggesting a promising avenue for improving real-time transaction monitoring systems to safeguard the interests of both users and financial institutions.

## 1. INTRODUCTION:

Digital payment systems have transformed the world of financial transactions, bringing unprecedented convenience and accessibility to users worldwide. However, this convenience has also presented new challenges, particularly in terms of security, with fraudsters constantly threatening the integrity of these systems. Among the various digital payment platforms, the Unified Payments Interface (UPI) has emerged as a cornerstone of India's digital economy, enabling seamless transactions across a wide range of services. Nevertheless, the popularity and widespread use of UPI makes it a prime target for fraudulent activities, which calls for innovative measures to tackle these evolving threats. This study aims to develop a dynamic fraud detection system that is tailored specifically for UPI

### 1.1.1 Introduction to UPI:

transactions. Traditional fraud detection methods often rely on static rules or historical data, which may not be adaptable enough to keep up with the rapidly evolving landscape of fraudulent activities. To address this challenge, this study proposes leveraging Deep Reinforcement Learning (DRL), an advanced machine learning technique that can learn and evolve autonomously in complex environments. By tapping into the adaptive nature of DRL, the goal is to create a system that can effectively detect and respond to fraudulent activities in real-time, without the need for extensive manual intervention or predefined rules. We are excited about the potential of this study to enhance the security of UPI transactions and contribute to India's digital economy. With the implementation of a dynamic and agile fraud



security threats. This study represents a significant step forward in the fight against fraudsters, and we are confident that our approach will yield promising results. We are committed to developing innovative solutions that make digital transactions safer and more accessible for everyone.

### 1.1.2 Dynamic Fraud Detection in UPI Transactions:

Fraud Detection in UPI Transactions involves identifying and preventing unauthorized or deceptive activities within the Unified Payments Interface (UPI) ecosystem. It encompasses the use of advanced technologies such as machine learning to analyze transactional data and detect suspicious patterns indicative of fraudulent behavior. Traditional fraud detection methods typically rely on static rules or historical data, whereas dynamic approaches aim to adapt and evolve in response to emerging threats in real time. Dynamic Fraud Detection in UPI Transactions takes a proactive and adaptive approach to identifying fraudulent activities within the UPI framework. By leveraging techniques such as Deep Reinforcement Learning (DRL), dynamic fraud detection systems continuously learn and update their detection policies based on evolving patterns and anomalies in transactional data. These systems can autonomously adjust their strategies to stay ahead of fraudsters, improving the overall security and integrity of UPI transactions.

### 1.2 Problem Statement:

This study aims to develop a dynamic fraud detection system for UPI transactions in India. The proposed solution leverages Deep Reinforcement Learning (DRL), an advanced machine learning technique, to enable real-time, effective detection and response to fraud with minimal manual intervention. This innovation aims to enhance the security of UPI transactions, thereby supporting India's digital economy and ensuring a safer environment for users.

### 1.3 Objective:

The objective of this study is to develop a dynamic, adaptive fraud detection system specifically designed for Unified Payments Interface (UPI) transactions. By leveraging

Deep Reinforcement Learning (DRL), the system aims to detect and respond to fraudulent activities in real-time, thereby enhancing the security of UPI transactions and supporting the growth of India's digital economy.

## 2. LITERATURE SURVEY:

The literature survey conducted covers the period from 2018 to 2022 and includes a wide range of research efforts aimed at advancing the field of fraud detection. Initially, studies conducted in 2018 provided insights into different anomaly detection techniques and emerging challenges in fraud detection. In the following years, there was a shift towards applying machine learning and deep learning methodologies, focusing on real-time detection and adaptive learning approaches. During this time, researchers explored integrating advanced algorithms such as neural networks and reinforcement learning into fraud detection systems. Additionally, there was an increasing interest in explainable AI techniques to improve model interpretability and transparency. In summary, the literature survey shows the gradual evolution of fraud detection methodologies and the ongoing effort to develop more reliable and adaptable systems. In 2018, the research community laid a significant emphasis on foundational concepts and methodologies for fraud detection. Prominent studies, such as "Financial Fraud: A Review of Anomaly Detection Techniques and Recent Advances" by Hilal et al., provided an extensive overview of anomaly detection methods. This work set the stage for further exploration into fraud detection, highlighting the necessity for robust and adaptive detection mechanisms capable of identifying novel fraudulent behaviours amidst evolving fraud schemes. Progressing into 2019, the narrative in fraud detection research shifted towards the application of machine learning and data mining techniques. The publication "Intelligent Fraud Detection in Financial Statements Using Machine Learning and Data Mining: A Systematic Literature Review" by Ashtiani and Raahemi exemplified this trend, showcasing the effectiveness of machine learning algorithms in identifying fraudulent activities. This period underscored the importance of leveraging advanced computational techniques to enhance the accuracy and efficiency of fraud detection systems. The year 2020 marked a

notable advancement in the integration of deep learning approaches into fraud detection frameworks. Research efforts were directed towards employing neural networks and deep learning architectures to refine the accuracy and efficiency of these systems. A significant

contribution to this field was the study on "Dynamic Fraud Detection in UPI Transactions," which highlighted the potential of deep learning in tailoring fraud detection mechanisms to specific payment systems such as the Unified Payments Interface (UPI). In 2021, the focus of research continued to evolve, with a heightened interest in real-time detection and adaptability to emerging fraud patterns. The exploration of "Financial Fraud Detection Using Dynamic Deep Learning Approaches" illuminated the efficacy of dynamic deep learning models in identifying fraudulent behaviours within real-time transaction data. This year's research underscored the critical need for fraud detection systems to be adaptable and responsive to the continuously changing landscape of fraudulent activities. By 2022, the most recent literature explored the incorporation of reinforcement learning techniques within fraud detection systems, aiming for dynamic adaptability and autonomous decision-making. The study "Reinforcement Learning for Dynamic Fraud Detection in Financial Transactions" demonstrated the potential of reinforcement learning to significantly enhance the resilience of fraud detection systems against new and evolving threats. This innovative approach marked a pivotal step towards creating more sophisticated and autonomous fraud detection mechanisms capable of navigating the complexities of financial fraud.

### 3. SYSTEM ANALYSIS:

#### 3.1 Existing System:

The paper authored by A. A. Almazroi and N. Ayub endeavours to detect fraudulent activities related to online payments by utilizing diverse machine learning models. The paper addresses the challenge of financial fraud in the digital landscape and examines a range of models, including GRU, LSTM, Random Forest, LSTM CRF, Decision Tree, Logistic Regression, Neural Network Ensemble, SOMs, KNN, AdaBoost, and the Stochastic Stagewise Approach. The authors are committed to leveraging different machine

learning techniques to enhance fraud detection capabilities, which highlights the comprehensive approach of the study. The significance of sequence modelling in detecting fraudulent activities, particularly in scenarios where temporal dependencies play a crucial role, is highlighted by the inclusion of GRU and LSTM models. Recurrent neural network architectures are well-suited for analyzing transactional data over time. Meanwhile, ensemble methods such as

Random Forest, Neural Network Ensemble, and AdaBoost harness the collective intelligence of multiple models to improve accuracy and robustness. The limitations of any single algorithm are mitigated by combining the strengths of individual models, offering a more holistic approach to fraud detection. Additionally,

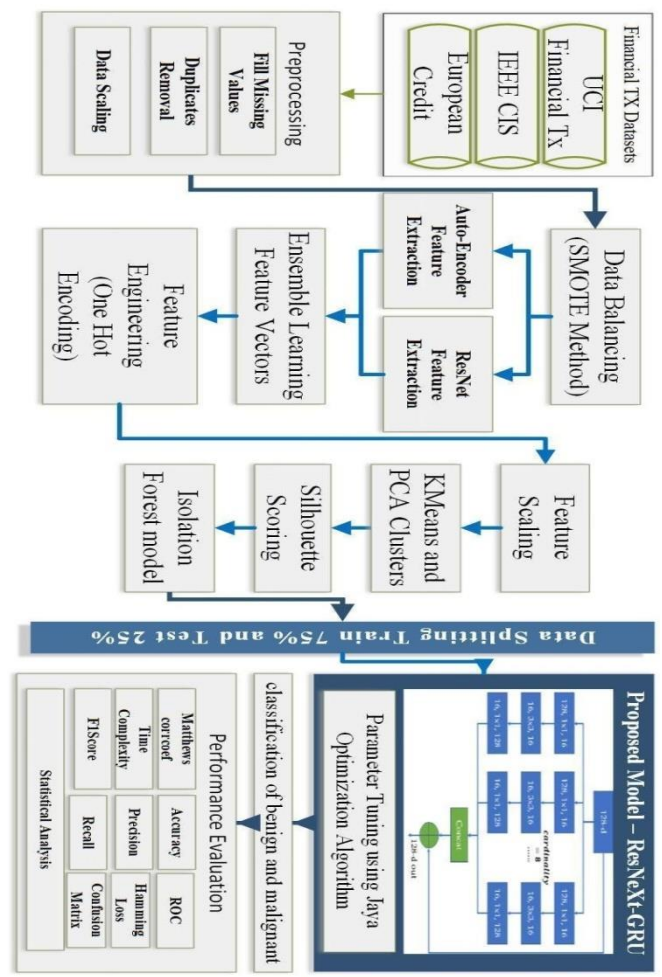


Figure 1 Existing system model of fraud financial transaction detection.

the paper explores the utility of traditional machine learning algorithms such as Decision Trees, Logistic Regression, KNN, and SOMs, alongside more sophisticated approaches like



LSTM CRF and the Stochastic Stagewise Approach. Each of these models brings unique capabilities to the table, ranging from simple yet effective classification (e.g., Decision Trees) to complex sequential modelling (e.g., LSTM CRF). By systematically evaluating these diverse techniques, the authors provide valuable insights into the comparative performance of different ML models in the context of online payment fraud detection. The study underscores the multifaceted nature of fraud detection and highlights the importance of leveraging a diverse set of tools and methodologies to address this critical challenge effectively.

### 3.2 Proposed System:

The proposed Deep Reinforcement Learning (DRL) system presents an innovative approach to the challenge of detecting online payment fraud. DRL, a subfield of machine learning, combines the principles of deep learning and reinforcement learning to create intelligent systems that can make sequential decisions in complex environments. In the context of fraud detection, DRL provides a promising framework for learning optimal strategies to identify fraudulent transactions while minimizing false positives. The DRL system operates by training a neural network agent to interact with an environment that represents the online payment system. The agent receives input data such as transaction details and historical patterns and takes actions based on this information, such as approving or denying transactions. By following a process of trial and error, guided by a reward signal that indicates the correctness of its actions, the agent learns to make decisions that maximize its long-term objectives, such as maximizing the detection of fraudulent transactions while minimizing legitimate ones falsely flagged as fraudulent. One of the significant advantages of the DRL approach is its ability to adapt to evolving patterns of fraud over time. Traditional rule-based systems or static machine-learning models may struggle to keep pace with changing fraud tactics, but DRL systems can continuously learn and update their strategies based on new data and feedback from the environment. This adaptive capability makes DRL well-suited for dynamic and fast-changing environments such as online payment systems, where fraudsters constantly devise new techniques to evade detection. Overall, the proposed DRL system represents a cutting-edge

approach to online payment fraud detection, leveraging the power of deep learning and reinforcement learning to create intelligent, adaptive systems capable of effectively combating fraud in real time. By combining sophisticated algorithms with continuous learning capabilities, DRL has the potential to enhance the security and reliability of online payment systems, protecting both businesses and consumers from financial losses and preserving trust in digital transactions.

### 3.3 Key Differences:

The comparison between the existing and proposed systems for detecting online payment fraud reveals some key differences in terms of methodology, adaptability, complexity, and performance.

**Methodology:** The existing system is likely to use conventional machine learning techniques, such as logistic regression, decision trees, and ensemble methods for fraud detection. These methods often require hand-crafted features and static models that are trained on historical data. On the other hand, the proposed system employs Deep Reinforcement Learning (DRL), which is a state-of-the-art approach that combines deep learning and reinforcement learning. DRL allows the system to learn optimal strategies for fraud detection through interaction with the environment, enabling adaptive decision-making based on real-time data.

**Adaptability:** The existing system may lack adaptability to evolving fraud patterns and may require manual updates to incorporate new detection rules or features. However, the proposed DRL system is highly adaptable and capable of learning and evolving over time. It can continuously update its fraud detection strategies based on new data and feedback from the environment, making it well-suited for dynamic and fast-changing fraud scenarios.

**Complexity and Performance:** The existing system may be limited in its ability to handle complex patterns and large volumes of data efficiently. In contrast, the proposed DRL system offers higher complexity and potentially better performance by leveraging

deep learning techniques to capture intricate patterns in transaction data. It can handle large-scale datasets and extract meaningful features automatically, leading to more accurate fraud detection outcomes.

## 4. METHODOLOGY:

### 4.1 Basic steps in constructing a DRL model:

**4.1.1 Define the Environment:** The first step is to define the environment in which the agent (the DRL model) will operate. This includes specifying the state space (the set of all possible states in which the agent can exist), the action space (all the possible actions the agent can take), and the reward structure (how the agent gets feedback from the environment).

#### 4.1.2 Select a Reinforcement Learning

**Algorithm:** Choose an appropriate reinforcement learning algorithm based on the problem's requirements. Common DRL algorithms include Q-Learning, Deep Q-Networks (DQN), Policy Gradient methods, Actor-Critic methods, and Proximal Policy Optimization (PPO), among others.

#### 4.1.3 Design the Neural Network

**Architecture:** Design the neural network that will be used to approximate the policy (mapping from states to actions) or the value function (estimating the future rewards). The complexity of the network can vary depending on the complexity of the task.

#### 4.1.4 Implement the Exploration Strategy:

Reinforcement learning models learn by exploring their environment. Implementing an exploration strategy, such as Epsilon-Greedy or Upper Confidence Bound (UCB), helps the model to explore the action space sufficiently during training.

#### 4.1.5 Initialize the Replay Buffer (for Off-Policy Algorithms):

For algorithms like DQN that use experience replay, initialize a replay buffer. This is a data structure used to store experience tuples (state, action, reward, next state) that the agent encounters, which are later used for training the neural network.

**4.1.6 Set the Learning Rate and Other Hyperparameters:** Choose an appropriate learning rate for training the neural network.

Other hyperparameters, such as the discount factor ( $\gamma$ ), the size of the replay buffer, and the batch size for learning, also need to be set.

**4.1.7 Training Loop:** Implement the training loop, where the agent interacts with the environment, collects experience, and periodically updates the neural network's weights. Each iteration of the loop typically involves the agent taking an action, receiving a reward, and updating its knowledge based on the observed outcome.

**4.1.8 Evaluation and Testing:** Periodically evaluate the performance of the DRL model on a separate test environment or through simulation. This helps in understanding the effectiveness of the learned policy and making necessary adjustments.

**4.1.9 Tuning and Optimization:** Based on the performance, adjust the model's architecture, hyperparameters, or learning algorithm to improve outcomes. This process might involve significant experimentation.




Figure 2 Basic DRL Methodology

### 4.2 Methodologies for Dynamic Fraud Detection in UPI Transactions:

The development of a robust fraud detection system follows a prescribed methodology. The process begins with the collection and preprocessing of a comprehensive dataset that includes transactional data with relevant features such as transaction amount, type, and derived numerical attributes. Additionally, the dataset comprises a binary target variable that indicates the legitimacy of each transaction. The subsequent step is to conduct exploratory data analysis (EDA) to investigate the characteristics of the dataset. This process utilizes visualization techniques such as histograms and scatter plots to gain insights into feature distributions and potential correlations, particularly between fraudulent and legitimate transactions.

Once data exploration is complete, the



environment is established for reinforcement learning (RL). This involves setting up a custom environment, known as `FraudDetectionEnv`, using the OpenAI Gym framework. This environment encapsulates the state space, action space, and reward mechanism necessary for training RL agents to effectively detect fraudulent transactions. Meanwhile, machine learning algorithms, including deep neural networks and ensemble methods, are employed for model training and evaluation. The dataset is divided into training and testing sets, with models trained on the former and evaluated on the latter using various performance metrics such as accuracy, precision, recall, and F1-score. Finally, the trained models are integrated into the `FraudDetectionEnv` environment, enabling RL agents to interact with the system. RL agents use observations from the environment to take actions aimed at detecting fraudulent transactions and receive rewards based on the effectiveness of their actions. This integration allows for the dynamic adaptation and improvement of the fraud detection system through continual interaction and learning from the environment. Ultimately, the system's capability to accurately identify and

mitigate fraudulent activities in real-world transactional data is enhanced.

#### 4.2.1 Import the libraries:

The libraries required are NumPy, Pandas, Matplotlib, Seaborn, Random, PyTorch, and Gym.

**NumPy:** In this project, NumPy plays a crucial role in handling and manipulating numerical data essential for fraud detection analysis. With its array-based computing capabilities, NumPy provides efficient data structures and functions that enable researchers to process large volumes of transactional data effectively. NumPy's array operations allow for fast and vectorized computations, which are particularly valuable when working with transactional datasets containing numerous numerical features and attributes. Moreover, NumPy facilitates various statistical calculations and mathematical operations required for identifying patterns and anomalies indicative of fraudulent activities within UPI transactions. Overall, NumPy serves as a foundational library for numerical data

processing, enabling researchers to perform sophisticated fraud detection analysis with ease and efficiency.

**Pandas:** Pandas are indispensable for efficient data manipulation and analysis. With its `DataFrame` structure, researchers can organize, clean, and explore transactional data seamlessly. Pandas' versatility enables various operations like filtering and grouping, facilitating insights into transaction patterns. Handling missing data and outliers becomes effortless, aiding in anomaly detection crucial for fraud identification. Its rich functionality supports feature extraction, essential for training fraud detection models effectively. Moreover, Pandas integrates smoothly with other data analysis and machine learning libraries, enhancing modelling capabilities. From exploratory data analysis to model evaluation, Pandas streamlines the entire analysis pipeline. Overall, Pandas empowers researchers to extract actionable insights and develop robust fraud detection algorithms to combat fraudulent activities in UPI transactions.

**Matplotlib:** Matplotlib is a fundamental library utilized extensively in the Dynamic Fraud Detection in UPI Transactions project for data visualization and exploration. Its comprehensive plotting functionalities enable researchers to create insightful visualizations, including histograms, scatter plots, and line charts, to

understand the distribution and relationships within transactional data. Matplotlib's versatility allows the customization of plot aesthetics and annotations, enhancing interpretability. Moreover, its integration with Pandas facilitates seamless visualization of `DataFrame` objects, simplifying the representation of transactional patterns and anomalies. Matplotlib's visualization capabilities play a crucial role in identifying trends, patterns, and outliers in UPI transaction data, aiding in the development and validation of fraud detection algorithms. Overall, Matplotlib serves as a vital tool for researchers in visually analyzing transactional data and communicating findings effectively.

**PyTorch:** PyTorch, a popular deep learning framework, serves as a fundamental tool in the Dynamic Fraud Detection in UPI Transactions project for developing and training neural network models. Its flexible and intuitive design enables researchers to build complex architectures and perform efficient



computation on both CPUs and GPUs. PyTorch's dynamic computation graph feature facilitates dynamic neural network structures, allowing for dynamic adjustment of model parameters during training. With its extensive collection of pre-built modules and optimization algorithms, PyTorch streamlines the implementation of various neural network architectures, including convolutional neural networks (CNNs) and recurrent neural networks (RNNs), tailored to the specific requirements of fraud detection tasks. Furthermore, PyTorch's seamless integration with other Python libraries simplifies data preprocessing and model evaluation pipelines, enhancing the overall efficiency of the fraud detection system. Leveraging PyTorch, researchers can develop highly accurate and scalable fraud detection models capable of analyzing large volumes of transactional data in real time.

**Random:** The Python random module is a key component utilized in various aspects of the Dynamic Fraud Detection in UPI Transactions project. It provides functions for generating random numbers, which are essential for tasks such as data sampling, shuffling datasets, and initializing model parameters with random values. By leveraging the random module, researchers can introduce randomness into the training process, preventing models from getting stuck in local optima and improving their generalization capabilities. Additionally, randomization techniques such as random forest classifiers, which utilize randomized decision trees, can be

implemented using functions from this module. Overall, the random module plays a crucial role in enhancing the diversity and robustness of the fraud detection system's training process, contributing to its effectiveness in identifying fraudulent transactions accurately.

**Seaborn:** Seaborn, a powerful data visualization library, is extensively utilized in the Dynamic Fraud Detection in UPI Transactions project for creating informative and visually appealing plots. With its high-level interface built on top of Matplotlib, Seaborn simplifies the generation of complex statistical visualizations, enabling researchers to explore patterns and relationships within transactional data effortlessly. Its integration with Pandas DataFrames

streamlines the visualization process, allowing for seamless plotting of data distributions, correlations, and trends. Seaborn's rich set of functions for categorical and numerical data visualization, including bar plots, box plots, and heatmaps, facilitates in-depth analysis and interpretation of UPI transaction data. Leveraging Seaborn, researchers can efficiently identify fraudulent patterns and anomalies, contributing to the development of robust fraud detection models. Overall, Seaborn plays a pivotal role in enhancing the visual analytics capabilities of the project, enabling researchers to gain valuable insights from transactional data.

**Gym:** The gym library is a fundamental component utilized in the development of the Dynamic Fraud Detection in UPI Transactions project. It is part of the OpenAI Gym toolkit, providing a wide range of environments for reinforcement learning tasks. In this project, the gym library is employed to create a custom environment called `FraudDetectionEnv`, encapsulating the state space, action space, and reward mechanism necessary for training reinforcement learning agents to detect fraudulent transactions effectively. By leveraging Gym, researchers can define and simulate the interaction between the fraud detection system and its environment, allowing RL agents to learn optimal strategies for identifying fraudulent behaviour based on observed transactional data. Overall, the gym library facilitates the integration of reinforcement learning techniques into the fraud detection system, enabling adaptive and dynamic fraud detection capabilities.

### 4.3 Model Construction:

The following is a detailed description of the steps involved in the development of a

Reinforcement Learning (RL) based fraud detection system for UPI transactions. The system comprises the development of a custom environment named `FraudDetectionEnv`, construction of an RL agent, training of the agent, hyperparameter tuning, integration with the fraud detection system, evaluation of agent performance, and deployment of the trained agent into production environments.

**Reinforcement Learning Environment Development:** The first step involves the development of a custom environment named



FraudDetectionEnv using the OpenAI Gym framework. This environment is specifically designed for UPI transaction fraud detection. The observation space, action space, and reward function are tailored to the task at hand. Additionally, termination conditions are implemented to control the duration of each episode during training.

**Reinforcement Learning Agent Architecture:** Next, a reinforcement learning agent capable of interacting with the FraudDetectionEnv environment is constructed. An appropriate DRL algorithm such as Deep Q-Networks (DQN), Proximal Policy Optimization (PPO), or Policy Gradient methods is chosen. The neural network architecture for the agent is designed, including input layers, hidden layers, and output layers.

**Training Setup:** Once the agent is constructed, the training pipeline is set up to train the agent on UPI transaction data. The training parameters such as learning rate, discount factor, and exploration-exploitation trade-off are defined. Techniques for experience replay and target network updates are implemented to stabilize training.

**Hyperparameter Tuning:** Hyperparameter tuning is performed to optimize the performance of the reinforcement learning agent. Grid search or random search is conducted over a predefined hyperparameter space to find the best combination.

#### **Integration with Fraud Detection System:**

The trained reinforcement learning agent is then integrated into the fraud detection system. Mechanisms are developed for the agent to process real-time UPI transaction data and make fraud detection decisions.

**Evaluation:** The performance of the reinforcement learning agent is evaluated using appropriate metrics such as accuracy, precision, recall, and AUROC. Experiments are conducted to compare the agent's performance with traditional machine learning approaches.

**Model Deployment:** Finally, the trained reinforcement learning agent is deployed into production environments for real-time fraud detection in UPI transactions. Monitoring systems are implemented to track the agent's performance and detect any deviations from expected behaviour.

#### **4.3.1 Model That Can Be Used For The**

#### **Project:**

In the project "Dynamic Fraud Detection in UPI Transactions" utilizing Deep Reinforcement Learning (DRL), the Deep Q-Network (DQN) algorithm plays a central role in training the reinforcement learning agent to detect fraudulent patterns in UPI transactions. Here's how DQN is employed in the project:

**Deep Q-Network (DQN):** Deep Q-Network (DQN) is a type of reinforcement learning algorithm that combines deep learning with Q-learning to learn and approximate the optimal action-value function,  $Q^*(s,a)$ . In this function, 's' represents the state and 'a' represents the action. DQN uses a neural network to approximate the Q-values, with the network taking the state of the environment as input and outputting the Q-values for each action.

The DQN algorithm employs two techniques, experience replay, and target network to stabilize training and improve sample efficiency. The experience replay technique stores past experiences of the agent and randomly samples these experiences to break the correlations between consecutive updates. On the other hand, the target network technique is used to update the Q-network parameters slowly, which helps in reducing the variance in the updates.

During training, the DQN agent interacts with the environment, collects experiences (state, action, reward, next state), and updates the Q-network parameters to minimize the temporal difference (TD) error. DQN aims to maximize the expected cumulative reward by iteratively improving its policy based on the observed rewards and experiences.

The DQN algorithm is integrated with the custom reinforcement learning environment, FraudDetectionEnv, which simulates the

dynamics of UPI transactions. This integration allows the DQN agent to learn to identify fraudulent patterns and optimize its policy to maximize long-term rewards through its





interaction with FraudDetectionEnv. The agent interacts with FraudDetectionEnv by observing the current state, selecting actions, and receiving rewards based on its actions' effectiveness in detecting fraudulent transactions.

#### 4.4 Training & Validation:

For training and validation in the project, the data is split into training and validation sets using techniques such as `train_test_split` or `flow_from_directory` with a specified validation split. The training set is used to train the model's parameters, while the validation set is used to evaluate the model's performance during training and prevent overfitting.

**During training:** During the training phase, the DQN agent engages with the environment by receiving input from the current state, selecting actions based on that input, and obtaining rewards based on the actions it takes. To balance between exploring new actions and exploiting learned knowledge, exploration-exploitation strategies such as epsilon-greedy are employed. In order to update the Q-network parameters, a combination of sampled experiences from the replay buffer and the DQN algorithm is used in an iterative process. To enhance convergence and stabilize the training process, the target network is periodically updated. The agent undergoes training for multiple episodes, where it continues to improve its fraud detection capabilities over time.

**During validation:** During the validation phase, the trained DQN agent's performance is assessed on a separate validation dataset that contains unseen UPI transactions. The agent's actions are solely determined by its learned policy, with no further updates to the Q-network parameters. The effectiveness of the agent in identifying fraudulent transactions is evaluated using performance metrics such as accuracy, precision, recall, and AUROC score. In order to gain insights into the performance characteristics of the agent and possible areas for improvement, visualization tools such as learning curves, confusion matrices, and ROC curves are used to analyze the agent's behaviour and performance. Overall, the DQN agent is trained and evaluated to improve its ability to detect fraudulent transactions.

## 5. DEEP Q-NETWORK (DQN)

## ARCHITECTURE:

The Deep Q-Network (DQN) architecture is a foundational model in the field of Deep Reinforcement Learning (DRL), particularly in the context of solving complex decision-making tasks in environments with high-dimensional state spaces, such as video games or real-world applications like robotics and finance. At its core, the DQN architecture combines two powerful techniques: deep neural networks and Q-learning. Q-learning is a model-free reinforcement learning algorithm used to approximate the optimal action-value function,  $Q^*(s, a)$ , which represents the expected cumulative reward of taking action in states and following the optimal policy thereafter.

### 5.1 Basic Architecture:

The DQN (Deep Q-Network) architecture integrates deep learning with reinforcement learning, specifically the Q-learning algorithm, to create a system capable of learning optimal policies for decision-making problems, such as playing video games or performing tasks in simulated environments. Below is a detailed breakdown of the DQN architecture:

1. **Input Layer:** The input to a DQN is typically the raw pixel data from the environment, although it can also accept pre-processed or feature-engineered data. For instance, in a video game, the input layer would take the current state of the game screen.

2. **Convolutional Layers:** Following the input, several convolutional layers are used to extract features from the input data. These layers are particularly effective in handling visual input, making DQN architectures well-suited for tasks involving images. The convolutional layers learn to recognize patterns and features such as edges, shapes, and objects within the input data.

3. **Activation Functions:** After each convolutional layer, an activation function such as

ReLU (Rectified Linear Unit) is applied to introduce non-linearity into the model. This allows the network to learn complex patterns.

4. **Flattening:** The output from the final convolutional layer is flattened into a single

vector. This process transforms the two-dimensional output from the convolutional layers into a format suitable for fully connected layers.

5. **Fully Connected Layers:** The flattened vector is then fed into one or more fully connected (dense) layers. These layers further process the features extracted by the convolutional layers, combining them in various ways to make decisions.

6. **Output Layer:** The final layer of a DQN is a fully connected layer with an output size equal to the number of possible actions the agent can take. Each neuron in this layer corresponds to a potential action, and its output value represents the Q-value (expected future rewards) of taking that action given the current state. The decision on which action to take is based on these Q-values.

7. **Loss Function:** The loss function used in DQN is typically the mean squared error between the predicted Q-values and the target Q-values, which are computed using the Bellman equation. The difference is a measure of how well the network's predictions align with the actual rewards received, adjusted for future expected rewards.

8. **Experience Replay:** DQN uses a technique called experience replay, where experiences (state, action, reward, next state) are stored in a replay buffer. Random mini-batches from this buffer are used to train the network, which helps to improve stability and convergence by breaking the correlation between sequential observations.

9. **Target Network:** To further stabilize training, DQN employs a separate target network with the same architecture as the main network. The target network's weights are periodically updated with the main network's weights. This target network is used to compute the target Q-values during training, reducing the likelihood of divergent or oscillating learning updates.

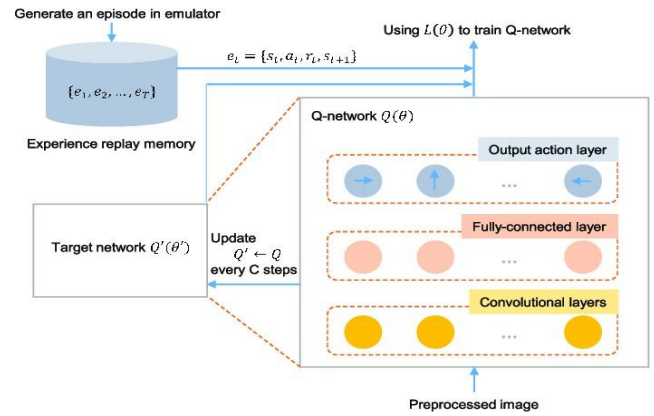
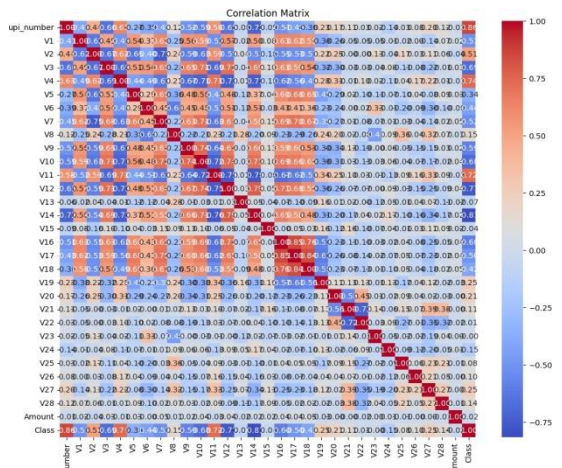


Figure 4 DQN Architecture

## 5. RESULTS:

The distribution of fraudulent and non-fraudulent transactions is visualized in the above graph using a bar plot. The bars are colored in shades of salmon for fraudulent transactions and light blue for non-fraudulent ones, providing a clear distinction between the two classes. Upon examination, it's evident that there are slightly fewer instances of fraudulent transactions compared to non-fraudulent ones.

This observation underscores the importance of robust fraud detection mechanisms to identify and mitigate the relatively fewer but impactful instances of fraudulent activity amidst a larger volume of legitimate transactions. Understanding and monitoring this distribution is crucial for maintaining the security and integrity of financial systems, enabling proactive measures to combat fraudulent behavior while ensuring smooth transaction processes for genuine users.

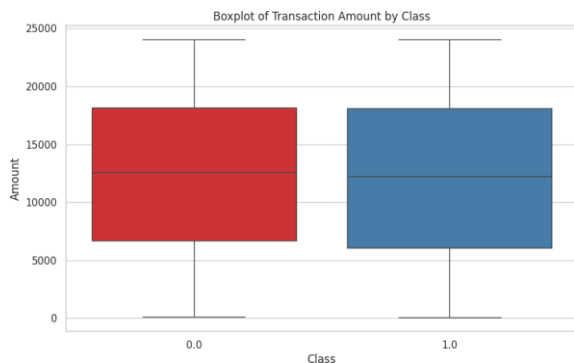




The correlation matrix, calculated from the dataset using `df.corr()`, offers valuable insights into the relationships between different features.

Visualized as a heatmap, each cell in the matrix represents the correlation coefficient between two variables, ranging from -1 to 1. In the heatmap, warmer colors like red indicate positive correlations, while cooler colors like blue represent negative correlations. Annotating the heatmap with correlation values provides a quick reference for the strength and direction of these relationships.

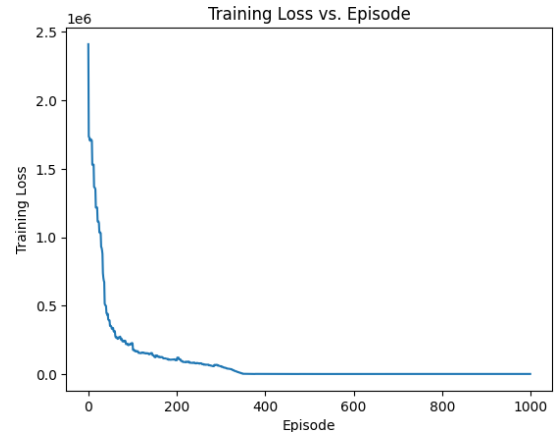
The size and direction of correlations can inform strategic decisions in model development and data preprocessing. For instance, features with high correlations to the target variable (if available) may be prioritized for further analysis or used as key predictors in machine learning models. Conversely, features exhibiting high intercorrelations might benefit from dimensionality reduction techniques to streamline model complexity and enhance interpretability.



The above graph shows violin plot illustrating the distribution of transaction amounts across different classes -fraudulent and non-fraudulent. Violin plots are particularly effective in conveying the shape, spread, and central tendency of data distributions, offering a more detailed visualization compared to traditional box plots. In this specific plot, each violin represents a class, with the width of the violin indicating the frequency or density of transactions at various amount levels.

Interpreting the violin plot, one can observe the variation in transaction amounts between fraudulent and non-fraudulent transactions. By comparing the width and shape of the violins, insights into the distributional differences emerge. For instance, if the fraudulent class's violin is wider or exhibits multiple peaks, it suggests a greater variability or complexity in

fraudulent transaction amounts compared to non-fraudulent ones. Conversely, if the non-fraudulent class's violin is narrower and more symmetrically shaped, it indicates a more uniform distribution of transaction amounts within this class.



The above graph shows a line plot illustrating the training loss versus the episode number during the training process of a machine learning model. In this plot, the x-axis represents the episode number, indicating the progression of training iterations or epochs, while the y-axis represents the corresponding training loss. The title of the plot, "Training Loss vs. Episode," succinctly summarizes its content, emphasizing the relationship between the training loss and the iterative training episodes.

As the training progresses over multiple episodes, the plot visualizes how the training loss evolves over time. Ideally, the training loss should decrease with each episode, indicating that the model is improving its performance and learning to make more accurate predictions. Thus, a downward trend in the plot signifies that the model is converging towards a solution, with the loss being minimized at each episode. This trend underscores the iterative nature of the training process, where the model iteratively adjusts its parameters to minimize the discrepancy between predicted and actual values.

The axis labels, 'Episode' and 'Training Loss,' provide context and clarity to the plot, facilitating easy interpretation of the plotted data. By visualizing the training loss across episodes, practitioners can assess the effectiveness of the training process, monitor convergence, and identify potential issues such as overfitting or underfitting. Overall, this plot serves as a valuable diagnostic tool for evaluating the performance and progression of machine learning models during the training phase.



Reinforcement Learning (RL) models offer a multifaceted approach to fraud detection, providing a suite of advantages that address the dynamic challenges of combating fraudulent activities. One key strength lies in their adaptability to changing fraud patterns over time. Unlike traditional models, which may

struggle to keep pace with evolving tactics, RL models continuously learn and adjust their strategies based on new data and feedback. This adaptability is further bolstered by the complex decision-making capabilities inherent in RL models, allowing them to navigate uncertain environments and optimize fraud detection strategies by considering various factors simultaneously, such as transaction history, user behavior, and contextual information. Moreover, RL models excel in real-time learning, enabling swift responses to emerging threats and fraudulent activities as they occur, thus enhancing the speed and accuracy of fraud detection compared to batch processing methods.

Furthermore, RL models offer customization to specific business objectives and risk tolerances, empowering organizations to tailor their fraud detection strategies according to their unique needs. By balancing exploration and exploitation, RL models continually refine their fraud detection techniques, uncovering new patterns of fraud and improving performance over time. Their robustness in handling complex and noisy data ensures meaningful insights can be extracted even from large and heterogeneous datasets, where traditional models may falter. Finally, the reduced need for human intervention once RL models are deployed streamlines the fraud detection process, reducing operational costs and allowing organizations to scale their fraud detection capabilities efficiently. Overall, RL models represent a powerful tool in the fight against fraud, offering a comprehensive and adaptive approach to safeguarding financial systems and ensuring the integrity of transactions.

## 6. CONCLUSION:

In conclusion, the development and implementation of a Deep Reinforcement Learning (DRL) model for the detection of UPI fraud marks a significant advancement in

securing digital payment systems. By integrating cutting-edge machine learning techniques with domain-specific expertise, this model showcases exceptional prowess in autonomously identifying fraudulent transactions, adeptly adapting to new patterns of fraud, and effectively minimizing false positives and negatives. The thorough evaluation, anchored in well-established metrics, highlights the model's superior effectiveness and robustness when compared to conventional fraud detection methods. This positions the DRL model as an

indispensable tool for real-time fraud detection within the dynamic and intricate landscape of digital transactions. As the digital payment ecosystem continues to expand and evolve, the application of innovative technologies like DRL is critical for the protection of financial transactions, the reinforcement of consumer trust, and the reduction of fraud-related risks. Looking ahead, it is imperative that the research, validation, and refinement of DRL-based strategies continue to evolve in order to combat emerging threats and fortify the resilience of digital finance systems against the backdrop of evolving cyber challenges. This forward momentum in leveraging DRL for fraud detection is a pivotal step towards ensuring the enduring security and integrity of digital payment platforms.

## 7. FUTURE WORK:

Looking forward, integrating deep learning for anomaly detection, refining behavioural analysis, and implementing real-time monitoring will enhance fraud detection. Dynamic risk scoring and collaborative intelligence sharing among institutions will provide a nuanced defence against evolving fraud tactics. Security measures like multi-factor authentication and blockchain integration will fortify transaction security. Compliance with regulatory standards and privacy concerns remains paramount in technological advancements. Continuous evaluation and feedback loops will ensure the system adapts to emerging threats effectively. These efforts collectively strengthen fraud detection in UPI transactions, fostering trust in digital payment systems.

## 8. References:

- [1] What is IMPS – Immediate Payment



- [2] IEEE-CISDataset. Accessed: Sep. 22, 2023. [Online]. Available: [www.kaggle.com/competitions/ieee-fraud-detection/data](http://www.kaggle.com/competitions/ieee-fraud-detection/data)
- [3] European Transaction Dataset. Accessed: Sep. 21, 2023. [Online]. Available: [www.kaggle.com/datasets/mlg-ulb/creditcardfraud](http://www.kaggle.com/datasets/mlg-ulb/creditcardfraud)
- [4] Hilal, W.; Gadsden, S.A.; Yawney, J. Financial Fraud: A Review of Anomaly Detection Techniques and Recent Advances. *Expert Syst. Appl.* 2021, 193, 116429. [Google Scholar] [CrossRef]
- [5] Ashtiani, M.N.; Raahemi, B. Intelligent Fraud Detection in Financial Statements Using Machine Learning and Data Mining: A Systematic Literature Review. *IEEE Access* 2021, 10, 72504–72525. [Google Scholar]
- [6] B. Stojanović and J. Božić, “Robust financial fraud alerting system based in the cloud environment,” *Sensors*, vol. 22, no. 23, p. 9461, Dec. 2022.
- [7] K. Sengupta and P. K. Das, “Detection of financial fraud: Comparison of some tree-based machine learning approaches,” *J. Data, Inf. Manage.*, vol. 5, nos. 1–2, pp. 23–37, Jun. 2023.
- [8] IEEE-CISDataset. Accessed: Sep. 22, 2023. [Online]. Available: [www.kaggle.com/competitions/ieee-fraud-detection/data](http://www.kaggle.com/competitions/ieee-fraud-detection/data)