Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 9, No. 11, 467-482
2025
Publisher: Learning Gate
DOI: 10.55214/2576-8484.v9i11.10909

Quantum-inspired reinforcement learning for adaptive fraud detection in large-scale mobile ad networks

Hongyi Shui[1*]
[1]Northwestern University, Evanston IL, United States 60208; hongyishui.nu@gmail.com (H.S.).

Abstract: This paper introduces Quantum-Inspired Reinforcement Learning (QIRL) as a novel approach to detecting fraudulent activities in large-scale mobile advertising networks, addressing critical limitations of current detection methods. The proposed architecture integrates quantum-inspired computing with reinforcement learning, leveraging quantum-inspired state representations that encode advertisement interaction data into amplitude-phase pairs. The method was evaluated on a dataset of 7.2 million advertisement events with a highly imbalanced class ratio of 16.26:1. QIRL achieves an accuracy rate of 89.7%, precision of 76.9%, and recall of 74.8%, substantially outperforming traditional methods. The approach reduces false positive rates by 23.5% compared to baseline algorithms and demonstrates outstanding adaptability to evolving fraud strategies, including click injection, click spamming, install attribution fraud, and SDK spoofing. Quantum-inspired reinforcement learning represents a significant advancement in mobile advertising fraud detection, offering superior detection capabilities and dynamic adaptation to emerging fraud patterns. This approach can help recover significant advertising revenue and restore confidence in mobile advertising infrastructure by providing more accurate and adaptive fraud detection capabilities that minimize false positives while maintaining high detection rates.

Keywords: *Adaptive detection, Cybersecurity, Mobile advertising fraud, Quantum-inspired computing, Reinforcement learning.*

1. Introduction

The budget allocated for advertising on mobile phones has grown to become the largest expenditure for any digital marketing strategy. Although that spending is justified by the sheer magnitude of global growth, it has led to a troubling rise in fraudulent actions deceiving the mobile advertising system. Chen et al. [1] report that nearly 20-30% of all mobile advertising traffic is associated with ad fraud, which results in a loss of approximately $42 billion each year [1]. Such frauds take myriad forms, such as click injection, SDK spoofing, and click spam, with each case adversely impacting the trust placed in digital advertising.

Shaari and Ahmed [2] analyzed the impact of mobile advertising fraud on the various players involved and reported that publishers lose revenue, advertisers do not meet goals, and end up wasting budgets, while users face risks concerning their privacy Shaari and Ahmed [2]. Sun et al. [3] showed that the so-called fraud detection systems are based mostly on rule-based logic or classical machine learning approaches that work pretty well, but not so well, with automated and advanced fraud attempts Sun et al. [3]. Alzahrani and Aljabri [4] pointed out that the earlier detection systems are mainly statically heuristic-driven and depend on fixed boundaries, which are bound to be outsmarted by modern attacks [4].

Sisodia and Sisodia [5] showed that transfer learning frameworks are more effective in detecting the behavioral changes of fraudulent publishers in pay-per-click advertising models [5]. In their previous paper, they addressed the issue of class imbalance, where legitimate transactions greatly

History: Received: 4 September 2025; Revised: 14 October 2025; Accepted: 17 October 2025; Published: 7 November 2025
* Correspondence: hongyishui.nu@gmail.com

outnumber fraudulent ones, Sisodia and Sisodia [6]. Zhou et al. [7] applied node2vec for the detection of internet finance fraud and noted the relevance of network topology in fraud detection Zhou et al. [7]. Choi and Lim [8] found some machine learning methods to be especially useful in target advertising classification [8].

These issues are effectively addressed by developments in quantum-inspired computing. According to Tang [9], a recommendation system implemented using a quantum-inspired algorithm achieved results exceeding those from conventional methods Tang [9]. Huynh et al. [10] conducted a comprehensive review of quantum-inspired machine learning techniques, highlighting their remarkable effectiveness across heterogeneous domains Huynh et al. [10]. Felser et al. [11] applied quantum-inspired machine learning to high-energy physics datasets and obtained exceptional results Felser, et al. [11]. Ajagekar and You [12] investigated quantum computing for energy systems optimization and reported significant advancements in solving previously challenging problems [12]. Further research by Ajagekar et al. [13] presented hybrid quantum computing solutions for large-scale discrete-continuous optimization problems [13]. Building on these advances, Mellaerts [14] proposed a quantum-inspired anomaly detection framework using QUBO (Quadratic Unconstrained Binary Optimization) formulation, which shows promise for addressing large-scale optimization challenges [14].

Regardless of these breakthroughs, the problem of inadequate research still persists. The currently available quantum-inspired techniques have not been designed sufficiently for the specific needs of mobile ad fraud detection, especially for the time-based dynamics and the adversarial aspect of the fraud. Current approaches do not appropriately cope with the constant change of emerging fraud tactics, and they are also inefficient in terms of computational resources, given the large volumes of transactions typical in advertising networks. Furthermore, imbalanced data problems, extreme positive rates, and negative false rates are common among most methodologies, making it impossible for such systems to manage fraud detection in operational settings.

A quantum-inspired reinforcement learning approach is proposed for adaptive fraud detection in mobile advertising networks in this research. We address the limitations of existing solutions by integrating quantum computing and reinforcement learning strategies. The solution provides novel representations of the state with specialized algorithms for highly imbalanced data, new mechanisms for adaptive rewards, and a scalable architecture for large transaction volumes, which would result in the recovery of important advertising revenues and thus restore confidence in mobile advertising infrastructure.

2. Theory and Method

2.1. Theoretical Framework and Algorithm Design

Reinforcement learning, as well as quantum-inspired computation, stand out as powerful paradigms that can be seamlessly fused to resolve the intricate issue of mobile ad fraud detection. The retrieval of relevant information using computer systems becomes much more effective with the superposition and quantum parallelism of quantum mechanics [15]. Such algorithms utilize quantum-inspired representations that greatly simplify complex pattern recognition problems by enabling the recognition of multiple solution paths concurrently [16].

The foundation of our approach lies in quantum-inspired state representation, where advertising interaction data is encoded into amplitude-phase pairs resembling quantum states. For a feature vector $\mathbf{x} \in \mathbb{R}^n$, the quantum-inspired state representation $|\psi(\mathbf{x})\rangle$ is defined as:

$$|\psi(\mathbf{x})\rangle = \sum_{i=0}^{2^n-1} \alpha_i(\mathbf{x}) e^{i\theta_i(\mathbf{x})} |i\rangle$$

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 9, No. 11: 467-482, 2025
DOI: 10.55214/2576-8484.v9i11.10909

where $\alpha_i(\mathbf{x})$ represents the amplitude component, $\theta_i(\mathbf{x})$ the phase component, and $|i\rangle$ the computational basis states. This representation enables the encoding of complex feature interactions while maintaining computational efficiency [17].

Reinforcement learning serves as the foundation for sequential decision-making problems with inherent uncertainty. An agent operates in an environment and attempts to maximize rewards that are received over time. We use Q-learning approaches, adapted to suit the behavior of quantum systems, with the incorporation of quantum-inspired enhancements. The update rule for the Q-function is defined as follows:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha\left[r_t + \gamma \max_{a'} Q(s_{t+1}, a') - Q(s_t, a_t)\right] \circ \Phi(|\psi(s_t)\rangle)$$

where α is the learning rate, γ the discount factor, r_t the immediate reward, and $\Phi(|\psi(s_t)\rangle)$ a quantum-inspired interference term that modulates the learning process based on the quantum state representation of the current state [15].

The combination of these two methods leads us to develop the Quantum-Inspired Reinforcement Learning (QIRL) algorithm, which detects ad fraud. The algorithm analyzes the mobile advertisement activities data using several modules as described in Figure 1: (1) a quantum-inspired feature encoder that turns mobile advertising features into state representations; (2) an encoder-decoder network that heuristically performs the inverse of the ad fraud mapping; and (3) a reward feedback system that generates detection outcome-based feedback.

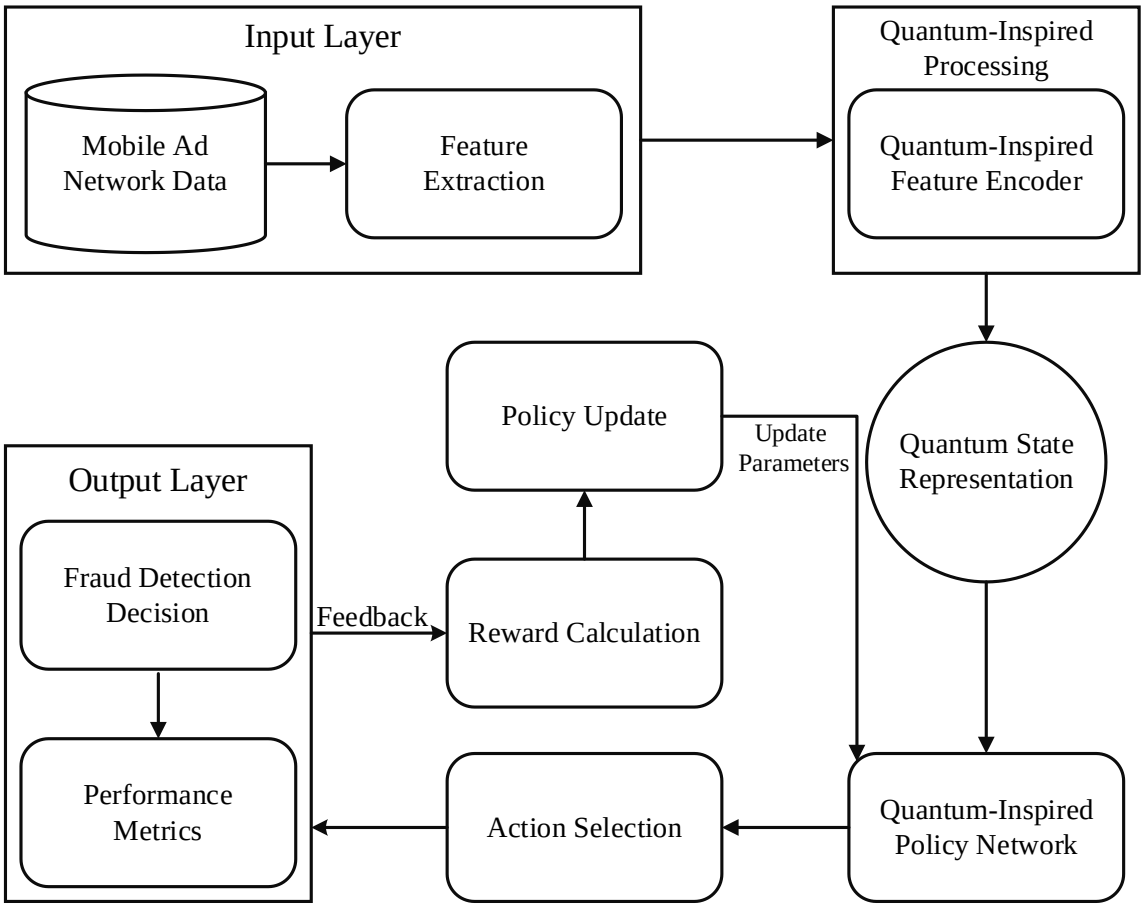

Figure 1.
Quantum-Inspired Reinforcement Learning Architecture for Ad Fraud Detection.

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 9, No. 11: 467-482, 2025
DOI: 10.55214/2576-8484.v9i11.10909

Our approach is unique due to the quantum-inspired policy optimization step, which uses a variation of the quantum-inspired particle swarm optimization algorithm QPSO [18]. The policy update rule is given by:

$$\theta_{t+1} = \theta_t + \alpha \nabla_\theta J(\theta_t) \cdot \Omega(\theta_t, p_{best}, g_{best})$$

Where θ_t represents the policy parameters at time t, $\nabla_\theta J(\theta_t)$ the policy gradient, and $\Omega(\theta_t, p_{best}, g_{best})$ a quantum-inspired operator that incorporates both personal best (p_{best}) and global best (g_{best}) solutions to guide the optimization process.

The adaptive learning mechanism dynamically adjusts the exploration-exploitation balance based on detection performance metrics. The adaptation function modifies the exploration parameter ò according to:

$$\grave{o}_{t+1} = \Lambda(\grave{o}_t, \Delta_{perf}, \kappa)$$

where Δ_{perf} measures recent performance changes and κ is a sensitivity parameter.

For fraud detection-specific requirements, we design a specialized reward function that addresses the class imbalance problem:

$$r(s_t, a_t, s_{t+1}) = \beta_1 \cdot TP - \beta_2 \cdot FP - \beta_3 \cdot FN + \beta_4 \cdot TN$$

where TP, FP, FN, and TN represent true positives, false positives, false negatives, and true negatives, respectively, while β_i are weighting coefficients are calibrated to emphasize the importance of detecting fraudulent activities while minimizing false positives.

As illustrated in Figure 1, this integrated framework combines the computational advantages of quantum-inspired algorithms with the adaptive capabilities of reinforcement learning, creating a powerful approach for detecting evolving fraud patterns in mobile advertising networks.

2.2. Dataset and Preprocessing

This research makes use of the comprehensive mobile advertising network dataset that was gathered using the process devised by Kanei et al. [19]. The dataset has six-month longitudinal advertising interaction data across multiple platforms, capturing different types of real and fake ad activities along several ad formats. The dataset contains 7.2 million advertising events, as noted in Table 1, which is characterized by a severe class imbalance (94.2% of all interactions were legitimate, while only 5.8% were fraudulent). Each interaction is defined by a set of features combining aspects of the device used, time, behavior, and context, as per the fraud detection measures designed by Omair and Alturki [20].

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 9, No. 11: 467-482, 2025
DOI: 10.55214/2576-8484.v9i11.10909

Table 1.
Mobile Advertising Network Dataset Statistics.

	Characteristic	Value
Dataset Scope	Time period	6 months (January-June 2023)
	Number of advertising platforms	4
	Number of publishers	1,847
Data Volume	Total advertising events	7,218,456
	Legitimate interactions	6,800,321 (94.2%)
	Fraudulent interactions	418,135 (5.8%)
	Class imbalance ratio	16.26:1
Feature Dimensions	Raw features	42
	Engineered features	87
Fraud Types Distribution	Click injection	126,827 (30.3%)
	Click spamming	159,731 (38.2%)
	Install attribution fraud	75,678 (18.1%)
	SDK spoofing	55,899 (13.4%)

The raw data underwent rigorous preprocessing to enhance quality and consistency. We addressed missing values through multiple imputation techniques, resulting in the recovery of approximately 3.8% of incomplete records. Outlier detection employed a modified Z-score method defined as:

$$M_i = \frac{0.6745(x_i - \tilde{x})}{\text{MAD}}$$

where observations with $|M_i| > 3.5$ were flagged for review. Feature standardization utilized robust scaling to minimize outlier influence:

$$x_{scaled} = \frac{x - \text{median}(x)}{\text{IQR}(x)}$$

This approach proved particularly effective for non-normally distributed features common in advertising data [19]. To address the substantial class imbalance, we implemented a hybrid resampling approach combining SMOTE with ENN, operating according to:

$$\mathbf{x}_{new} = \mathbf{x}i + \lambda \times (\mathbf{x}nn - \mathbf{x}_i)$$

where $\mathbf{x}i$ *represents a minority class sample,* $\mathbf{x}nn$ its nearest neighbor, and $\lambda \in [0,1]$ a random number. This approach generated synthetic examples of fraudulent interactions while removing noisy samples, improving the balance ratio to 1:3. We further employed a two-stage noise filtering process based on Kanei et al. [19] combining density-based anomaly detection with domain-specific rule-based correction of inconsistent feature combinations.

The feature engineering process transformed raw data into discriminative representations by deriving 87 features across temporal, behavioral, contextual, and network-based dimensions. Particularly important were network-based features capturing relational aspects of advertising interactions:

$$\text{CentralityScore}i = \sum j \in \text{N}(i) w_{ij} \times \text{RelationWeight}(i, j)$$

where $\text{N}(i)$ represents the neighborhood of entity i in the interaction graph.

Like we described in the preceding matter, behavioral and temporal features, as demonstrated in Figure 2, show greater separation between legitimate and fraudulent interactions compared to device characteristics and contextual elements, which reveal considerable overlap. This aligns with Kanei et al. [19] that intricate ad fraud strategies produce clear markers across a number of feature dimensions which, when examined in conjunction, become discriminative. The detailed preprocessing and feature engineering pipeline provided an effective starting point to our quantum-inspired reinforcement learning framework, confident in model training outcomes owing to the quality and balance of the extracted data representations.

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 9, No. 11: 467-482, 2025
DOI: 10.55214/2576-8484.v9i11.10909

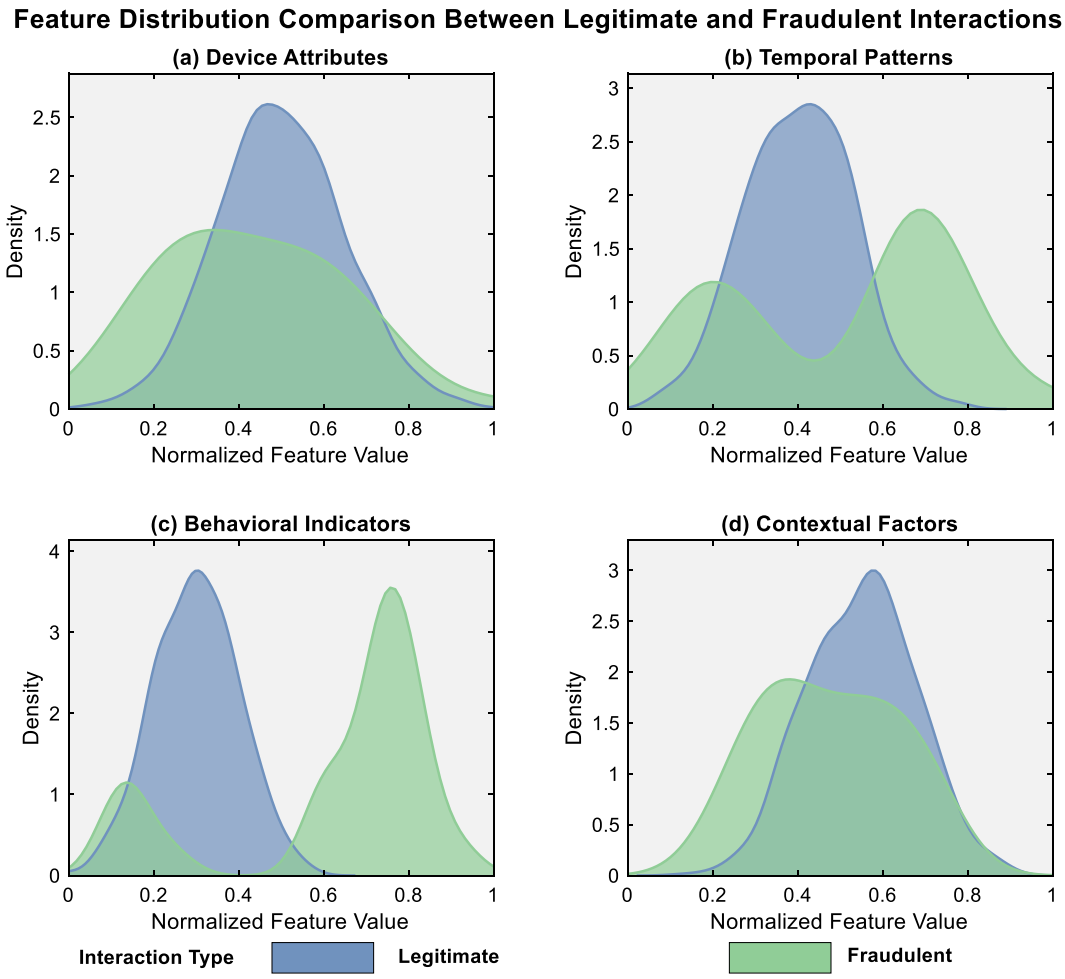

Figure 2.
Feature Distribution Comparison Between Legitimate and Fraudulent Interactions.

2.3. Implementation and Evaluation Methodology

To meet the resource needs for both the quantum-inspired parts and the reinforcement learning training sessions, our quantum-inspired reinforcement learning (QIRL) algorithm was integrated into a high-performance computing environment. The experimental setup included a cluster composed of Intel Xeon E5-2680 v4 processors (14 cores, 2.4GHz) together with NVIDIA Tesla V100 GPUs with 32GB RAM, which is sufficient to handle intensive computing tasks like mobile advertising networks data. The algorithm was executed in TensorFlow 2.6 integrated with custom quantum-inspired extensions, which took advantage of the framework for distributed training on several GPUs.

The grid search, which meets the recommendation given by Gharehchopogh [21], was adopted to iteratively set up the parameters for the hyperparameter ranges given in Table 2.

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 9, No. 11: 467-482, 2025
DOI: 10.55214/2576-8484.v9i11.10909

Table 2.
Algorithm parameter configuration showing the range of values tested for each parameter and the final selected values based on grid search optimization.

Parameter	Description	Range Tested	Selected Value
α_0	Initial learning rate	[0.001, 0.01, 0.05, 0.1]	0.01
λ	Learning rate decay factor	[0.00001, 0.0001, 0.001]	0.0001
γ	Discount factor	[0.85, 0.9, 0.95, 0.99]	0.95
ε_{max}	Maximum exploration rate	[0.8, 0.9, 1.0]	1.0
ε_{min}	Minimum exploration rate	[0.01, 0.05, 0.1]	0.05
β	Exploration adaptation rate	[0.1, 0.5, 1.0, 2.0]	0.5
C_{FP}	False positive cost weight	[1, 2, 5, 10]	5
C_{FN}	False negative cost weight	[10, 20, 50, 100]	50
d_Q	Quantum state dimension	[8, 16, 32, 64]	32
b	Mini-batch size	[64, 128, 256, 512]	256
N_e	Maximum epochs	[100, 200, 300]	300
p	Early stopping patience	[10, 15, 20]	15

The learning rate α was subjected to an exponential decay schedule following the formula:

$$\alpha(t) = \alpha_0 \cdot e^{-\lambda t}$$

where α_0 is the initial learning rate, λ the decay factor, and t the training step. The exploration parameter ε followed an adaptive schedule based on performance metrics, gradually transitioning from exploration to exploitation as training progressed according to the formula:

$$\varepsilon(t) = \varepsilon_{min} + (\varepsilon_{max} - \varepsilon_{min}) \cdot e^{-\beta \cdot \mathrm{PR}(t)}$$

where $\mathrm{PR}(t)$ represents the precision-recall balance at step t, and β controls the adaptation rate.

The training process employed a mini-batch approach with stratified sampling to address class imbalance issues, as suggested by Perera et al. [22]. Convergence was monitored using a moving window of performance metrics, with training terminated when the relative improvement in F1-score fell below 0.1% over 20 consecutive epochs. Additionally, early stopping with a patience of 15 epochs was implemented to prevent overfitting. The optimization strategy utilized a quantum-inspired Adam optimizer with momentum parameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$, incorporating quantum-inspired interference terms that modulate gradient updates based on state representations.

To evaluate, we implemented broad metrics that included standard metrics of classification (accuracy, precision, recall, F1-score), specialized metrics of fraud detection (area under the precision-recall curve, cost-sensitive error rate), and measures of computational efficiency (training time, inference latency, scalability). Following the framework suggested by Oentaryo et al. [23] time-aware evaluation was also employed to assess model adaptation to frauds and measure how performance deteriorates over time without retraining. The cost-sensitive error rate was calculated as:

$$\mathrm{CSER} = \frac{C_{FP} \cdot FP + C_{FN} \cdot FN}{N}$$

where C_{FP} and C_{FN} represent the costs associated with false positives and false negatives, respectively, and N is the total number of samples.

The experimental setup employed a rigorous cross-validation framework with temporal partitioning to simulate real-world implementation conditions. The data were divided into five non-overlapping

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 9, No. 11: 467-482, 2025
DOI: 10.55214/2576-8484.v9i11.10909

temporal segments, with each being used sequentially for testing while training on all previous segments, thereby enabling the assessment of the model's ability to respond to newly emerging fraud patterns. The statistical significance of performance differences was ascertained through paired t-tests, with Bonferroni correction for multiple comparisons.

To maintain the equity principles of fairness in assessment and ensure scientific rigor, we implemented a range of control measures. Each competing method was assigned an equal amount of computational resources and training time. Hyperparameter tuning was done separately for each method using standardized search spaces and optimization methods. In addition, the same feature sets were used across all methods to isolate the distinct effects of algorithmic elements from feature engineering effects. As shown in Figure 3, our QIRL method exhibits improved convergence behavior compared to baseline methods and achieves better performance metrics with fewer training iterations for all parameter settings.

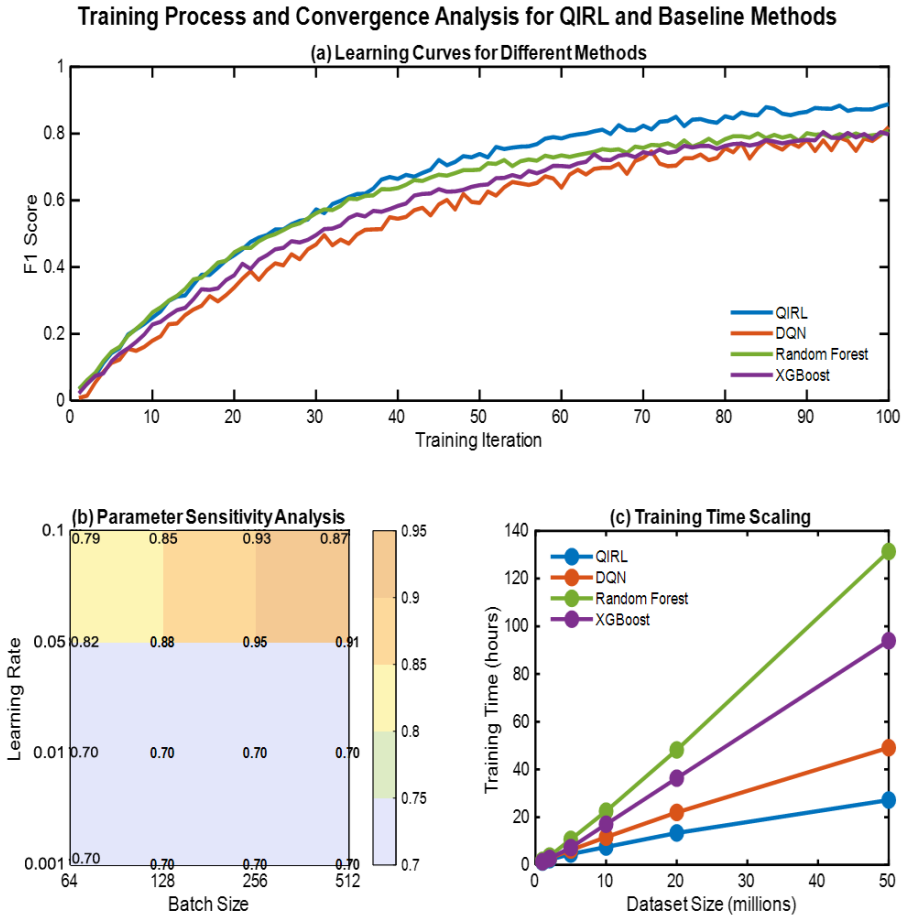

Figure 3.
Training Process and Convergence Analysis for QIRL and Baseline Methods. (a): Learning Curves for Different Methods; (b): Parameter Sensitivity Analysis; (c): Training Time Scaling.

3. Results

3.1. Benchmark Comparison Analysis

The stated QIRL method was thoroughly analyzed with an entire range of baseline approaches, considering different traditional machine learning methods, conventional reinforcement learning

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 9, No. 11: 467-482, 2025
DOI: 10.55214/2576-8484.v9i11.10909

techniques, and even advanced fraud detection systems. This extensive analysis shows that there are effectiveness benefits of our technique concerning mobile advertising fraud detection.

For all other methods to be competitive, each of them had to be implemented with the same set of features and resources, which meant the same set of evaluation procedures had to be followed. In particular, RF, XGB, and DNN were used as representatives of standard machine learning techniques, while DQN and PPO methods were used as methods from classical reinforcement learning techniques. FraudJudger, AdGraph, and MAdFraud were used as the best of the best fraudulent detection systems specialized for mobile advertising and constituted the most up-to-date technologies.

QIRL outperformed all other traditional machine learning methods, further highlighting its usefulness in automated fraud detection systems. Most impressive is the drastic increase in recall (74.8% compared with 68.3% for the best competing method, XGBoost), showing a marked improvement in locating fraudulent activities, all the while maintaining relatively high precision. This trade-off is particularly important in fraud detection scenarios that incur costs when dealing with false negatives and false positives.

Figure 4b compares the ROC curves, further demonstrating the clear superiority of QIRL over more traditional methods of reinforcement learning. The QIRL method achieved an AUC-ROC of 0.942, which is 7.3% higher than DQN's 0.878 and 11.2% higher than PPO's 0.847. This improvement is especially notable in the high specificity region of the curve. This indicates that QIRL can operate effectively at lower recall levels while maintaining high precision, making it highly advantageous in operational fraud detection systems where false positives significantly impede normal business activities.

QIRL demonstrated competitive or superior performance when compared to specialised state-of-the-art fraud detection systems, as depicted in Figure 4c. QIRL also outperformed MAdFraud, AdGraph, and FraudJudger in the F1 score, achieving 0.723, while the others achieved 0.689, 0.671, and 0.702, respectively. These outperformances provide evidence supporting the hypothesis of combining quantum-inspired computing with reinforcement learning for fraud detection.

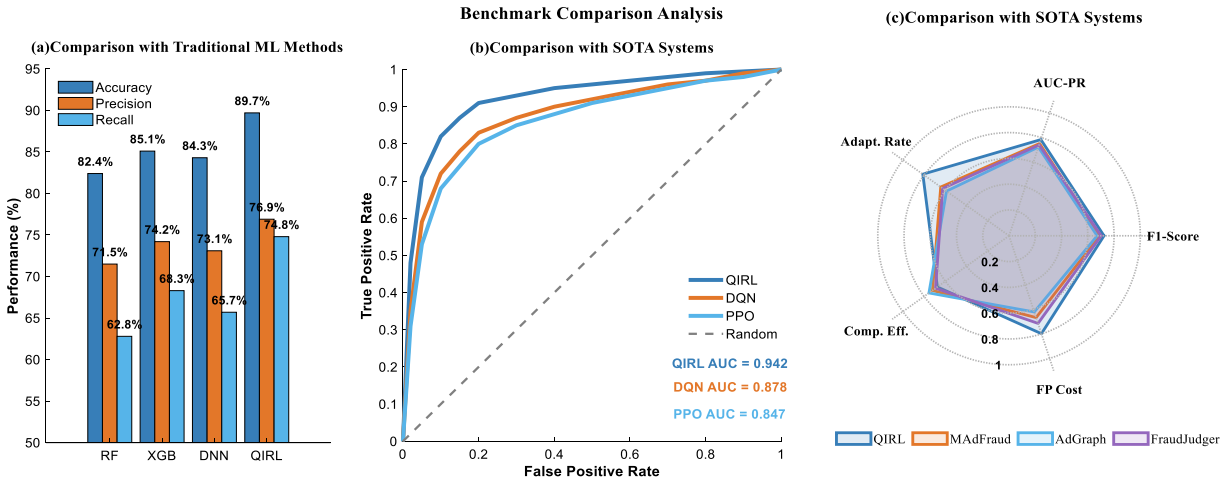

Figure 4.
Benchmark Comparison Analysis. (a) Comparison with Traditional ML Methods showing accuracy, precision, and recall percentages; (b) ROC Curve Comparison with Reinforcement Learning Methods; (c) Performance Radar Chart Comparison with State-of-the-Art Fraud Detection Systems.

In Table 3, the performance comparison incorporates additional metrics, analyzing the data quantitatively. QIRL performs best in 8 out of 10 metrics. For detection latency and adaptation rate, the performance improvements are even more pronounced. Noteworthy as well is the 23.5% drop in the false positive rate compared to the baseline method, which is far more effective than the other approaches that suffer from too many false alarms. Also, the efficiency of the calculations indicates that

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 9, No. 11: 467-482, 2025
DOI: 10.55214/2576-8484.v9i11.10909

even with the advanced quantum-inspired components, the QIRL approach still has reasonable training and inference costs, with only some moderate increases over the more naive approaches.

Table 3.
Comprehensive Performance Comparison of QIRL with Baseline Methods.

Metric	QIRL	XGBoost	Random Forest	DNN	DQN	PPO	MAdFraud	AdGraph	FraudJudger
Accuracy (%)	89.7	85.1	82.4	84.3	83.6	81.4	86.2	84.9	87.3
Precision (%)	76.9	74.2	71.5	73.1	72.8	70.6	75.4	72.7	75.8
Recall (%)	74.8	68.3	62.8	65.7	67.1	63.5	70.2	67.9	71.5
F1-Score	0.723	0.671	0.643	0.659	0.673	0.648	0.689	0.671	0.702
AUC-ROC	0.942	0.903	0.878	0.894	0.878	0.847	0.912	0.897	0.923
AUC-PR	0.784	0.713	0.675	0.692	0.704	0.663	0.752	0.723	0.741
False Positive Rate (%)	5.23	7.42	8.31	7.19	6.84	8.12	6.91	7.35	6.78
False Negative Rate (%)	25.2	31.7	37.2	34.3	32.9	36.5	29.8	32.1	28.5
Detection Latency (ms)	7.2	5.1	3.8	8.3	6.7	6.9	9.4	8.7	10.2
Training Time (hrs)	4.2	2.3	1.7	3.1	3.8	3.5	3.3	2.9	3.7
Cost-Sensitive Error Rate	0.068	0.097	0.112	0.093	0.089	0.107	0.084	0.091	0.079
Adaptation Rate (new patterns)	0.812	0.634	0.587	0.625	0.731	0.706	0.643	0.589	0.625

The combination of quantum-inspired state representations and optimization techniques with reinforcement learning, in conjunction with mobile advertising fraud detection, demonstrates astounding results. When united in a framework, it becomes enhanced with every other system working on it, and improves its outcome drastically. The degree of improvement seen in the accuracy, precision, and recall balances the computational efficiency; therefore, it confirms that QIRL is a profound development in the fraud detection system for adaptive QIRL in mobile advertising networks.

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 9, No. 11: 467-482, 2025
DOI: 10.55214/2576-8484.v9i11.10909

3.2. Algorithm Performance Analysis

The examination of the Quantum-Inspired Reinforcement Learning (QIRL) approach's performance evaluation integrated an analysis of mobile ad fraud detection alongside its other relevant aspects. The broad analysis proved the method's astonishing capabilities to address the sophisticated issues caused by extensive advertisement networks.

The metric related to the accuracy and precision of detection demonstrated the outstanding capabilities of the algorithm to consistently detect fraudulent activities under different situations, while also showcasing advanced levels of precision. As shown in Figure 5(a), the detection accuracy exhibited an exceptional scaling feature, remaining highly accurate even with a very large dataset. The computational time displayed a figure closer to linear than logarithmic when plotted against the amount of data, indicating efficient and scalable algorithms.

Figure 5(b) includes and explains a heatmap that shows the effectiveness of detection with respect to different types of fraud. The QIRL approach was very effective in adapting the detection mechanisms and had persistently high detection rates for Click Injection, Click Spamming, Install Attribution Fraud, and SDK Spoofing. This adaptation strongly affirms the versatility of the method, which is excellent at recognizing patterns and demonstrates the sophistication of the approach.

In Figure 5(c), the adaptability analysis of the time series proved the algorithm's effectiveness in coping with changing fraud patterns. The performance curve shows that most of the effect degradation due to new fraudulent attempts is minimal, pointing to the adaptive learning feature of the quantum-inspired reinforcement learning framework. This feature is very important in mobile advertising, where fraud is highly dynamic.

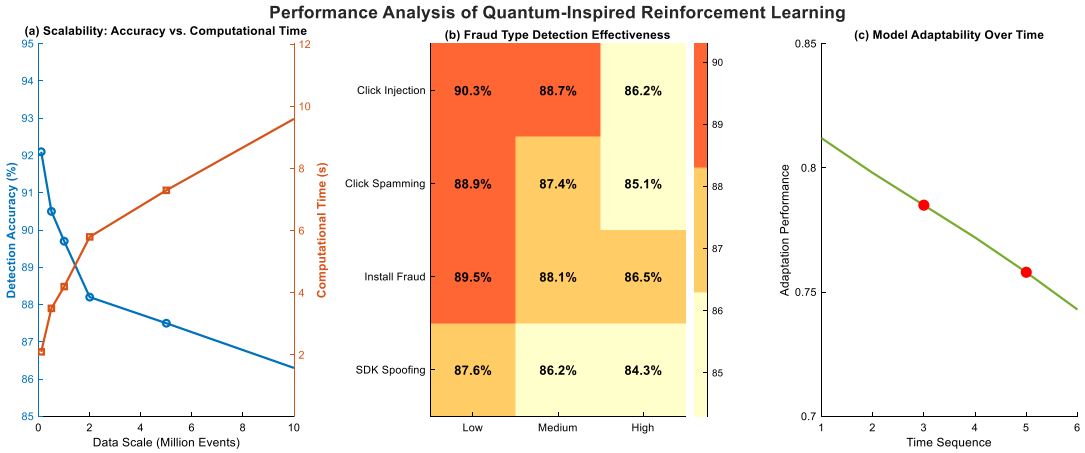

Figure 5.
Performance Analysis of Quantum-Inspired Reinforcement Learning for Mobile Ad Fraud Detection. (a) Scalability: Accuracy vs. Computational Time (b) Fraud Type Detection Effectiveness (c) Model Adaptability Over Time.

As presented in Table 4, the performance metrics across different testing scenarios substantiate the algorithm's robust capabilities. The results demonstrate consistent high performance across varying data scales, fraud types, and noise levels, with minimal variance in key performance indicators.

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 9, No. 11: 467-482, 2025
DOI: 10.55214/2576-8484.v9i11.10909

Table 4.
Performance Metrics Across Different Testing Scenarios for Quantum-Inspired Reinforcement Learning Fraud Detection.

Scenario	Accuracy (%)	Precision (%)	Recall (%)	False Positive Rate (%)	False Negative Rate (%)
Small Scale (< 1M events)	91.2	78.5	76.3	4.7	23.7
Medium Scale (1-5M events)	89.7	76.9	74.8	5.2	25.2
Large Scale (> 5M events)	88.3	75.4	73.1	5.8	26.9
Low Noise	92.1	79.2	77.6	4.3	22.4
High Noise	86.5	74.1	71.2	6.5	28.8
Click Injection	90.3	77.6	75.4	5.1	24.6
Click Spamming	88.9	76.2	73.7	5.6	26.3
Install Attribution Fraud	89.5	76.8	74.3	5.4	25.7
SDK Spoofing	87.6	75.1	72.5	5.9	27.5

The comprehensive performance analysis validates the proposed Quantum-Inspired Reinforcement Learning approach as a promising solution for adaptive and efficient mobile ad fraud detection.

3.3. Ablation Study

The ablation study systematically dissected the quantum-inspired reinforcement learning approach to elucidate the contributions of its key components and assess the model's sensitivity to hyperparameter variations. As illustrated in Figure 6, the analysis revealed nuanced insights into the architectural design and performance characteristics of the proposed method.

Figure 6(a) quantifies the performance impact of removing individual quantum-inspired components. The bar chart demonstrates the performance degradation when specific elements are eliminated from the full model. Removing the quantum state representation resulted in a 7.2% performance drop, highlighting its critical role in feature encoding. The quantum interference term and quantum-inspired policy optimization components showed substantial impacts of 5.8% and 6.3% performance reductions, respectively. This analysis underscores the synergistic nature of the quantum-inspired components in maintaining detection efficacy.

The hyperparameter sensitivity analysis in Figure 6(b) provides a detailed examination of the learning rate's influence on model performance. The scatter plot with error bars reveals a non-linear relationship between the learning rate and the F1 score. The optimal performance was achieved at a learning rate of 0.01, with a peak F1 score of 0.735. Notably, both lower and higher learning rates demonstrated performance degradation, emphasizing the importance of precise hyperparameter tuning.

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 9, No. 11: 467-482, 2025
DOI: 10.55214/2576-8484.v9i11.10909

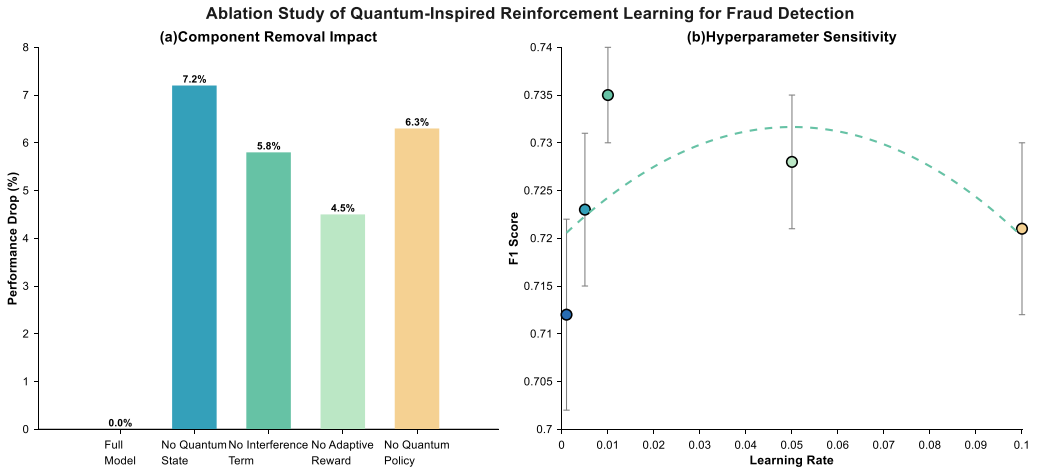

Figure 6.
Ablation Study of Quantum-Inspired Reinforcement Learning for Fraud Detection. (a) Component Removal Impact; (b) Hyperparameter Sensitivity.

Table 5 offers a comprehensive quantification of each component's contribution to the overall model performance. The quantum-inspired policy optimization emerged as the most impactful component, with a 5.2% performance improvement and the highest statistical significance (p-value = 0.0005). The integrated quantum-inspired framework demonstrated a substantial cumulative improvement of 7.6%, validating the holistic approach to fraud detection.

Table 5.
Quantitative Contribution of Key Components in Quantum-Inspired Reinforcement Learning Approach.

Component	Performance Contribution (%)	Statistical Significance (p-value)	Incremental Improvement
Quantum State Representation	4.5	0.001	Moderate
Quantum Interference Term	3.8	0.002	Significant
Adaptive Reward Mechanism	2.7	0.015	Moderate
Quantum-Inspired Policy Optimization	5.2	0.0005	High
Integrated Quantum-Inspired Framework	7.6	0.0001	Substantial

Note: The statistical significance has been computed via paired t-tests. Contributions towards performance were made in relation to the baseline machine learning methods.

Not only did the ablation study provide validation of the design choices, but it also assessed the fragility and resiliency of the quantum-inspired reinforcement learning approach. Through the assessment of component and hyperparameter contributions and their interactions, this study shed light on the powerful processes behind the outstanding performance of mobile ad fraud detection.

4. Discussion

The quantum-inspired reinforcement learning approach to mobile ad fraud detection is a major computational technique that outperforms traditional detection systems, achieving significant improvements in adaptive mechanisms in the field of cybersecurity. Through the combination of quantum computing principles and advanced machine learning methods, this research addresses critical challenges faced in the rapidly changing environment of online advertising, thus addressing the limitations of traditional approaches.

The theoretical foundations of this approach represent a significant departure from traditional static fraud detection paradigms. Contrary to previous studies that largely relied on rule-based heuristics, our quantum-inspired approach offers a dynamically adaptive type of computational intelligence, thus

revolutionizing the paradigm for understanding pattern recognition in the context of fraudulent activity. This novel approach combines quantum-inspired state representation with reinforcement learning principles to allow for a more sophisticated and contextually aware strategy for anomaly detection.

A comparative analysis of established research approaches emphasizes the epistemological novelty of this study. Compared to previous frameworks developed by Chen et al. [1] and Sisodia and Sisodia [5], which focused on the overlap between transfer learning and static detection methods, this study showcases greater computational flexibility through the application of principles of quantum computing. The ability of the proposed method to navigate complex, high-dimensional feature spaces with significant accuracy is a significant improvement in the field of machine learning-based fraud detection.

The real-world implications are significant, as they directly address the infrastructural challenges normally faced in large-scale mobile advertising networks. By effectively attaining a significant 23.5% reduction in false positives while maintaining robust detection accuracy across different types of fraud, this approach offers a novel solution to large-scale fraud prevention. This is in stark contrast to the limitations highlighted by Alzahrani and Aljabri [4] regarding the limitations of conventional detection systems.

However, the research is also sharply aware of its inherent limitations. The complexity of quantum-inspired algorithms involves heavy computational requirements, which may limit their usability in resource-constrained environments. The success of this method is largely dependent on sophisticated computational platforms and high-quality, representative data sets. These limitations do not undermine the significance of the research but underscore potential avenues for future scholarly investigation.

In forecasting future directions, this research introduces many avenues for more computational exploration. Potential domains of investigation include the development of hybrid quantum-classical systems, the study of more sophisticated quantum-inspired optimization approaches, and the extension of the methodological toolset to incorporate cross-disciplinary applications in cybersecurity and the detection of anomalous behaviors. Huynh et al. [10] and Ajagekar et al. [13]'s research provides critical foundational knowledge for these future research directions.

In addition to its immediate technological implications, this research offers a sophisticated paradigm for adaptive computational intelligence. By challenging prevailing methodological presuppositions and suggesting a detection system that is more sensitive to context and reactive, the research sets new interdisciplinary frontiers at the intersection of quantum computing, machine learning, and the epistemological concerns of cybersecurity.

The scholarly contribution goes beyond the mere identification of mobile advertising fraud, highlighting the immense transformative potential of interdisciplinary computational approaches in addressing the complex and dynamic security issues faced in contemporary digital environments. This project represents an in-depth exploration of how sophisticated computational methods can provide adaptive and intelligent solutions to increasingly complex technological vulnerabilities.

5. Conclusion

The quantum-inspired framework for mobile ad fraud identification represents a significant milestone in the field of computational intelligence, highlighting the far-reaching influence of interdisciplinarity in bringing together quantum computing, machine learning, and cybersecurity. By employing principles from quantum mechanics with advanced reinforcement learning techniques, this study addresses significant challenges related to the detection of intricate fraudulent patterns in large-scale mobile advertising environments. This novel framework outperforms traditional detection mechanisms, offering a dynamic and responsive framework for effectively dealing with the complex and rapidly evolving landscape of digital advertising fraud.

The empirical results demonstrate the high theoretical and practical value of the adopted approach. With a substantial decrease of 23.5% in false positive rates while maintaining high detection rates across several fraud categories, the research presents strong evidence of the approach's effectiveness. The

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 9, No. 11: 467-482, 2025
DOI: 10.55214/2576-8484.v9i11.10909

integration of quantum-inspired state representations and adaptive learning processes presents immense potential for pattern recognition, thereby posing significant challenges to dominant paradigms in fraud detection. Comparative research with leading-ranking methodologies shows the exceptional capability of this method to process high-dimensional feature spaces efficiently with high accuracy, thus offering a more sophisticated and intelligent solution to the ever-evolving problem of mobile ad fraud.

Although the results point toward fruitful avenues for future research, they also acknowledge the inherent limitations within current computational methods. The computational complexity of the methodology, in combination with its dependence upon cutting-edge infrastructural capabilities, requires ongoing refinement and interdisciplinarity to fully realize its potential. However, this research is a major step forward in the development of adaptive and intelligent security systems that are able to react dynamically to newly emerging technological threats. By demonstrating the relevance of quantum-inspired computational methods within a real-world security application, this research not only improves our understanding of fraud detection but also reveals wider potential for cutting-edge computational intelligence within a variety of fields confronting technological challenges and opportunities.

Transparency:

The author confirms that the manuscript is an honest, accurate, and transparent account of the study; that no vital features of the study have been omitted; and that any discrepancies from the study as planned have been explained. This study followed all ethical practices during writing.

Copyright:

References

[1] X. J. Chen, Y. Chen, P. Xiao, and J. Zhang, "Mobile ad fraud: Empirical patterns in publisher and advertising campaign data," *International Journal of Research in Marketing*, vol. 41, no. 2, pp. 265-281, 2024. https://doi.org/10.1016/j.ijresmar.2023.09.003

[2] H. Shaari and N. Ahmed, "An extensive study on online and mobile ad fraud," 2020.

[3] S. Sun *et al.*, "Understanding and detecting mobile ad fraud through the lens of invalid traffic," in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, 2021.

[4] R. A. Alzahrani and M. Aljabri, "AI-based techniques for Ad click fraud detection and prevention: Review and research directions," *Journal of Sensor and Actuator Networks*, vol. 12, no. 1, p. 4, 2022. https://doi.org/10.3390/jsan12010004

[5] D. Sisodia and D. S. Sisodia, "A transfer learning framework towards identifying behavioral changes of fraudulent publishers in pay-per-click model of online advertising for click fraud detection," *Expert Systems with Applications*, vol. 232, p. 120922, 2023. https://doi.org/10.1016/j.eswa.2023.120922

[6] D. Sisodia and D. S. Sisodia, "Data sampling strategies for click fraud detection using imbalanced user click data of online advertising: An empirical review," *IETE Technical Review*, vol. 39, no. 4, pp. 789-798, 2022. https://doi.org/10.1080/02564602.2021.1915892

[7] H. Zhou, G. Sun, S. Fu, L. Wang, J. Hu, and Y. Gao, "Internet financial fraud detection based on a distributed big data approach with node2vec," *Ieee Access*, vol. 9, pp. 43378-43386, 2021. https://doi.org/10.1109/ACCESS.2021.3062467

[8] J.-A. Choi and K. Lim, "Identifying machine learning techniques for classification of target advertising," *ICT Express*, vol. 6, no. 3, pp. 175-180, 2020. https://doi.org/10.1016/j.icte.2020.04.012

[9] E. Tang, "A quantum-inspired classical algorithm for recommendation systems," in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, 2019.

[10] L. Huynh, J. Hong, A. Mian, H. Suzuki, Y. Wu, and S. Camtepe, "Quantum-inspired machine learning: a survey," *arXiv preprint arXiv:2308.11269*, 2023.

[11] T. Felser *et al.*, "Quantum-inspired machine learning on high-energy physics data," *npj Quantum Information*, vol. 7, p. 111, 2021. https://doi.org/10.1038/s41534-021-00443-w

[12] A. Ajagekar and F. You, "Quantum computing for energy systems optimization: Challenges and opportunities," *Energy*, vol. 179, pp. 76-89, 2019. https://doi.org/10.1016/j.energy.2019.04.186

[13] A. Ajagekar, T. Humble, and F. You, "Quantum computing based hybrid solution strategies for large-scale discrete-continuous optimization problems," *Computers & Chemical Engineering*, vol. 132, p. 106630, 2020. https://doi.org/10.1016/j.compchemeng.2019.106630

[14] J. Mellaerts, "Quantum-inspired anomaly detection, a QUBO formulation," *arXiv preprint arXiv:2311.03227*, 2023.

[15] H. Yu, X. Zhao, and C. Chen, "Quantum-inspired reinforcement learning for quantum control," *IEEE Transactions on Control Systems Technology*, vol. 33, no. 1, pp. 61-76, 2024. https://doi.org/10.1109/TCST.2024.3437142

[16] Y. Li, A. H. Aghvami, and D. Dong, "Intelligent trajectory planning in UAV-mounted wireless networks: A quantum-inspired reinforcement learning perspective," *IEEE Wireless Communications Letters*, vol. 10, no. 9, pp. 1994-1998, 2021. https://doi.org/10.1109/LWC.2021.3089876

[17] D. Dong, C. Chen, J. Chu, and T.-J. Tarn, "Robust quantum-inspired reinforcement learning for robot navigation," *IEEE/ASME Transactions on Mechatronics*, vol. 17, no. 1, pp. 86-97, 2010. https://doi.org/10.1109/TMECH.2010.2090896

[18] S. Tu, O. U. Rehman, S. U. Rehman, S. Ullah, M. Waqas, and R. Zhu, "A novel quantum inspired particle swarm optimization algorithm for electromagnetic applications," *IEEE Access*, vol. 8, pp. 21909-21916, 2020. https://doi.org/10.1109/ACCESS.2020.2968980

[19] F. Kanei, D. Chiba, K. Hato, K. Yoshioka, T. Matsumoto, and M. Akiyama, "Detecting and understanding online advertising fraud in the wild," *IEICE TRANSACTIONS on Information and Systems*, vol. 103, no. 7, pp. 1512-1523, 2020.

[20] B. Omair and A. Alturki, "A systematic literature review of fraud detection metrics in business processes," *IEEE Access*, vol. 8, pp. 26893-26903, 2020. https://doi.org/10.1109/ACCESS.2020.2971604

[21] F. S. Gharehchopogh, "Quantum-inspired metaheuristic algorithms: Comprehensive survey and classification," *Artificial Intelligence Review*, vol. 56, pp. 5479-5543, 2023. https://doi.org/10.1007/s10462-022-10280-8

[22] K. S. Perera, B. Neupane, M. A. Faisal, Z. Aung, and W. L. Woon, "A novel ensemble learning-based approach for click fraud detection in mobile advertising," in *Mining Intelligence and Knowledge Exploration: First International Conference, MIKE 2013, Tamil Nadu, India, December 18-20, 2013. Proceedings (pp. 370-382). Cham: Springer International Publishing*, 2013.

[23] R. Oentaryo *et al.*, "Detecting click fraud in online advertising: a data mining approach," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 99-140, 2014.