



DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: <https://www.svedbergopen.com/>



Research Paper

Open Access

Dynamic Portfolio Optimization Algorithm Using Deep Q Networks for Corporate Financial Decision Systems

Dr.M. Sangeetha^{1*}, Dr. Sapna Bawankar², Abdumutalliev Abdulakhad Abdusamad Ugli³, Maksuda Dusmetova⁴, Nigora Uralova⁵, Kattakul Kinjaev⁶

¹Professor, Department of Information Technology, K.S.Rangasamy College of Technology, Tiruchengode, India.

E-mail: sangeetham@ksrct.ac.in, <https://orcid.org/0000-0003-4648-4242>

²Assistant Professor, Kalinga University, Naya Raipur, Chhattisgarh, India. E-mail: ku.sapnabawankar@kalingauniversity.ac.in,

<https://orcid.org/0009-0005-9651-582X>

³Turan International University, Namangan, Uzbekistan. E-mail: aabdulaxadturan@gmail.com,

<https://orcid.org/0009-0006-1675-0756>

⁴Senior Teacher, Urgench State Pedagogical Institute Urgench, Uzbekistan. E-mail: maksudaxon.70@gmail.com, <https://orcid.org/0009-0009-8269-5312>

⁵PhD, Senior Lecturer, Department of Romance-Germanic Languages, Jizzakh State Pedagogical University, Jizzakh, Uzbekistan.

E-mail: nigoraurolova026@gmail.com, <https://orcid.org/0000-0002-9627-9469>

⁶Lecturer, Department of finance and tourism, Termez University of Economics and Service, Termez, Uzbekistan.

E-mail: samurai6356693@gmail.com, <https://orcid.org/0009-0002-9315-1395>

*Corresponding author: Email: sangeetham@ksrct.ac.in

Abstract

This paper introduces DQN-PO, a new Deep Q-Network Portfolio Optimization framework designed specifically for corporate financial systems dealing with the challenges of high frequency, multi-asset trading environments. Standard Mean-Variance Optimization suffers from non-stationarity and transaction costs issues; DQN-PO resolves these problems by formulating the process of portfolio rebalancing as an MDP with continuous actions. The novel method uses a state space composed of 64 features representing price momentum, technical measures, macroeconomic data, and relationships between assets within the portfolio. Overfitting is prevented by implementing dual networks, prioritized experience replay, and double Q-learning strategies. In tests on 10 S&P 500 sectors over 2019-2023 (covering the COVID-19 market turmoil period), DQN-PO shows a Sharpe ratio of 2.14 and a 38.6% annual return compared to Markowitz MVO (1.43 Sharpe), risk parity (1.31 Sharpe), and baseline reinforcement learning with PPO agent (1.72 Sharpe). The maximum drawdown is controlled at -8.3% against -14.2% for Markowitz MVO. Ablation experiments show a significant contribution of dueling architecture (+11% Sharpe), prioritized experience replay (+8%), and multi-factor states (+19%).

Keywords: Deep Q-Networks, Reinforcement Learning, Portfolio Optimization, Financial Decision Systems, Markov Decision Process, Sharpe Ratio Maximization, Dueling Network Architecture, Experience Replay.

This is an open access article under CC BY 4.0, allowing unrestricted use with proper attribution, a license link, and indication of any changes made.

1. Introduction

Modern corporate treasury and asset management functions face a serious challenge: capital allocation through a variety of financial instruments under realistic conditions such as transaction costs, leverage limitations, ESG constraints, and fast-changing market microstructure. Classical portfolio theory based on Markowitz's (1952) mean-variance optimization approach operates under the assumption of a constant probability distribution and quadratic utility, assumptions which do not hold during regime switches such as the one experienced in 2008 or in 2020 due to the coronavirus outbreak [1].

Reinforcement learning (RL) appears to be a suitable approach. In its turn, by treating the portfolio management task as a sequential decision process, an RL-based agent learns policies for optimal cumulative risk-adjusted returns without any need to make distributional assumptions [15] [21][23]. DQN, initially developed by Mnih et

al. (2015) for the Atari video game playing domain, proved to be capable of estimating Q-values over high-dimensional state spaces, an approach applicable in finance [2][15].

Even though there is much to be gained from using DQN for financial portfolios, there are some problems that make its implementation problematic. These include (1) continuous action space, which needs discretization or approximation, (2) non-stationary reward signals as a result of regime shifts, (3) over-estimation of Q-value in volatile markets, and (4) sparsity of reward signals during drawdown periods.

1.1 Research Contributions

The contributions of this paper can be summarized as follows:

1. A multi-factor 64-dimensional state representation including technical, fundamental, and macro indicators for a 10-asset universe.
2. Adapted Duelling Double DQN with Prioritized Experience Replay for portfolio optimization.
3. Custom reward function involving transaction costs, drawdowns, and Sharpe-ratio shaping.
4. Extensive empirical evaluation against seven baselines over 5 years of out-of-sample performance, including a market crash period.
5. Component-wise ablation study measuring the value of each individual architectural element.

1.2 Paper Organization

Section 2 provides an overview of prior work on portfolio management via reinforcement learning. Section 3 formalizes the MDP and its reward design. Section 4 introduces DQN-PO architecture. Section 5 outlines experiments performed. Section 6 shows the results obtained. Section 7 interprets them.

2. Related Work

2.1 Classical Portfolio Optimization

Markowitz's mean-variance optimization (MVO) from 1952 is a classic method of quantitative finance. Variations include the Black-Litterman model from 1990, which combines the beliefs of investors with equilibrium return levels, and the risk-parity method (Qian, 2005), which seeks to achieve equality of marginal risk contributions [3][7]. Factor models [6] [20] consider returns in terms of risk premia and provide insights into portfolio allocation decisions [6]. However, while elegant in theory, such methods presuppose stable covariance estimates, which is negated by the empirical observations of volatility clustering, fat tails, and structural breaks [10] [16].

2.2 Reinforcement Learning in Finance

The first uses of reinforcement learning algorithms to solve finance-related problems are seen in the studies conducted by Moody et al. (1998) [12]. In this case, recurrent policy networks were applied in foreign exchange operations. This research is validated by the data presented by Aripin et al. (2025) in the form of the Crypto Portfolio Management data set [4][18]. This data set proves that the use of policy gradient algorithms beats the Markowitz method in selecting portfolios in cryptocurrencies. Yang et al. (2020) used Deep Deterministic Policy Gradient (DDPG), an RL algorithm that works well with continuous action spaces in actor-critic frameworks in investing in US stock portfolios [5][22].

2.3 Deep Q-Networks for Discrete Financial Actions

The constraint posed by DQN of requiring discrete actions has led to various adaptation methods. Deng et al. (2016) have used discretization of portfolio weights to obtain a finite set of actions (underweight, neutral, overweight for each asset), which allowed them to update their Q-table using the tabular method. In Carta et al. (2021), double DQN with dueling heads was applied to a portfolio of 30 stocks and obtained a Sharpe ratio of 1.8 for Italian stock index funds [8][9][19]. The proposed DQN-PO uses all improvements introduced to the DQN method in a multi-dimensional state space.

3. Problem Formulation

3.1 Markov Decision Process Definition

Formalize portfolio rebalancing as an MDP defined by the tuple $M = (S, A, P, R, \gamma)$, where:

- $S \subseteq \mathbb{R}^{64}$ is the state space capturing market and portfolio features at each time step t .

- $A = \{a_1, \dots, a_K\}$ is the discrete action set representing $K = 2,187$ permissible weight configurations.
- $P: S \times A \times S \rightarrow [0,1]$ is the (unknown) transition kernel of the financial market.
- $R: S \times A \rightarrow \mathbb{R}$ is the reward function defined in Section 3.3.
- $\gamma = 0.97$ is the discount factor, balancing immediate versus future reward.

Time steps correspond to daily market closes. The agent observes s_t , selects action a_t (target allocation vector w_t), and receives reward r_t before transitioning to s_{t+1} .

3.2 State Space Design

The state vector s_t of 64 dimensions is classified into six groups of features, as presented in table 1 below. Feature engineering decisions are based on the market microstructure research [11][24], and tested via ablation (Section 6.3). All the features are scaled via z-score normalization within a 63-day look-back window, which corresponds to a quarterly period.

Table 1: State Space Decomposition — Feature Groups and Dimensionality for N=10 Asset Universe

State Component	Variables	Description	Dimensionality
Price Momentum	P_t, ΔP_t, ROC_20	Close, 1-day Δ, 20-day rate-of-change	3 × N_assets
Technical Indicators	RSI, MACD, BBands	Oscillators, trend, volatility bands	5 × N_assets
Macro Factors	VIX, Yield Curve, FX	Risk regime, rates, currency exposure	8 scalars
Portfolio State	w_t, PnL_t, DD_t	Current weights, P&L, drawdown metrics	N_assets + 4
Volume & Liquidity	OBV, ADR, Spread	On-balance vol, avg daily range, bid-ask	3 × N_assets
Sentiment Signals	NLP Score, ESG	News sentiment, sustainability ratings	2 × N_assets
TOTAL STATE DIM	—	For N=10 asset universe	64 scalars

Components of the 64-dimensional state vector s_t . BN = Batch Normalization; OBV = On-Balance Volume; ADR = Average Daily Range. NLP scores are computed via a pretrained FinBERT model applied to daily news headlines.

3.3 Action Space Discretization

Portfolio weight allocation on a continuous scale is considered using discrete weights allocated as either underweights {5%}, neutral weights {10%}, or overweights {15%} for each asset. In this case, the number of actions is $3^N=59,049$ when $N=10$. Budget and diversification restrictions limit the number of feasible actions to $K=2,187$ actions. All actions are pre-calculated and stored in a lookup table to ensure $O(1)$ look-up times at prediction time.

3.4 Reward Function

The reward at time t is a composite of three terms:

$$r_t = \alpha \cdot SR_t - \beta \cdot TC_t - \delta \cdot DD_t$$

where SR_t is the rolling 20-day Sharpe ratio computed from portfolio log-returns, $TC_t = \kappa \cdot \|w_t - w_{t-1}\|_1$ is the $L1$ transaction cost penalty with $\kappa = 10$ bps per unit turnover, and $DD_t = \max(0, peak_t - NAV_t)/peak_t$ penalizes active drawdown. Coefficients $\alpha = 1.0, \beta = 0.3, \delta = 0.5$ are tuned via cross-validated hyperparameter search.

4. DQN-PO Architecture

4.1 Network Design

DQN-PO employs a dueling double DQN (D3QN) architecture in table 2. The network splits into two streams after the shared representation layers: a value stream $V(s; \theta_V)$ estimating the expected return from state s , and an advantage stream $A(s, a; \theta_A)$ estimating the relative benefit of action a . The Q-value is reconstructed as:

$$Q(s, a; \theta) = V(s; \theta_V) + [A(s, a; \theta_A) - \max_{\{a\}} A(s, a'; \theta_A)]$$

This formulation enables the network to learn state value independently of action advantage, improving stability when many actions share similar values—common in portfolio rebalancing when market conditions are stable.

Table 2: DQN-PO Neural Network Architecture — Layer Specifications

Layer	Type	Input Dim	Output Dim	Activation	Parameters
Input Layer	Dense (State)	64 features	256 units	Linear	16,640
Hidden Layer 1	Dense + BN	256 units	512 units	ReLU	131,584
Hidden Layer 2	Dense + Dropout	512 units	512 units	ReLU	262,656
Hidden Layer 3	Dense + BN	512 units	256 units	ReLU	131,328
Hidden Layer 4	Dense	256 units	128 units	ReLU	32,896
Output Layer (Q)	Dense	128 units	N actions	Linear	$N \times 129$

Network architecture for DQN-PO. BN = Batch Normalization applied after linear transformation. Dropout ($p=0.20$) applied during training only. The output layer produces Q-values for all $K=2,187$ feasible portfolio allocation actions.

4.2 Training Components

4.2.1 Prioritized Experience Replay

Transitions (s_t, a_t, r_t, s_{t+1}) are stored in a sum-tree replay buffer $|D| = 100,000$. Sampling priority is proportional to $|\delta_i|^\alpha$ where δ_i is the TD error and $\alpha=0.6$. Importance sampling weights $w_i = (N \cdot P(i))^{-\beta}$ correct for the non-uniform distribution with β annealed from 0.4 to 1.0 over training, preventing over-weighting of high-error transitions during early learning.

4.2.2 Target Network and Double Q-Learning

A separate target network $Q^{\{-\}}(s, a; \theta^{\{-\}})$ is updated every $\tau_{freq} = 1,000$ steps by hard-copying the online network parameters. Double Q-learning decouples action selection (online network) from Q-value evaluation (target network), correcting the maximization bias inherent to standard DQN:

$$y_t = r_t + \gamma \cdot Q^{\{-\}}(s_{t+1}, \operatorname{argmax}_{\{a\}} Q(s_{t+1}, a; \theta); \theta^{\{-\}})$$

4.2.3 Epsilon-Greedy Exploration Schedule

During training, action selection follows an ϵ -greedy policy. ϵ decays linearly from $\epsilon_0 = 1.0$ to $\epsilon_{min} = 0.01$ over $T_\epsilon = 50,000$ steps. After decay, the agent maintains 1% random exploration to prevent policy collapse in low-volatility regimes.

4.3 Hyperparameter Configuration

Table 3 reports all hyperparameters, their selected values, search ranges, and search methods. Tuning was conducted via random search with 200 trials on the validation period (2018–2019), optimizing for Sharpe ratio.

Table 3: Hyperparameter Settings — DQN-PO Training Configuration

Hyperparameter	Symbol	Value	Search Range	Method
Learning Rate	α	3×10^{-4}	$[10^{-5}, 10^{-3}]$	Log-uniform
Discount Factor	γ	0.97	$[0.90, 0.999]$	Linear grid
Replay Buffer Size	$ D $	100,000	$[50K, 500K]$	Log-uniform
Mini-Batch Size	B	256	$[64, 512]$	Powers of 2
Target Net Update Freq	τ_{freq}	1,000 steps	$[500, 5000]$	Log-uniform
Epsilon Initial	ϵ_0	1.0	Fixed	—
Epsilon Final	ϵ_{min}	0.01	$[0.001, 0.1]$	Log-uniform
Epsilon Decay Steps	T_ϵ	50,000	$[20K, 200K]$	Log-uniform
Hidden Layer Size	H	512	$[128, 1024]$	Powers of 2
Dropout Rate	p_{drop}	0.20	$[0.10, 0.50]$	Linear grid
Transaction Cost (bps)	κ	10	Fixed (market)	—

Final hyperparameter configuration after random search over 200 trials on 2018–2019 validation data.

5. Experimental Setup

5.1 Data and Universe

Evaluate DQN-PO on daily OHLCV data for 10 SPDR sector ETFs from 01/01/2015 to 31/12/2023 sourced from Yahoo Finance. The asset universe covers XLK (Technology), XLV (Healthcare), XLF (Financials), XLE (Energy), XLC (Communication), XLY (Consumer Discretionary), XLP (Consumer Staples), XLRE (Real Estate), XLB (Materials), and XLU (Utilities). The nine-year span encompasses two complete market cycles and includes the disruption caused by COVID-19.

The data is divided chronologically: training (2015–2017), validation (2018–2019), and out-of-sample testing (2020–2023). All performance metrics provided are calculated solely based on the out-of-sample testing window to avoid lookahead bias [13].

5.2 Baseline Methods

For benchmarking DQN-PO, select six approaches from the realms of classic optimization and reinforcement learning:

- Markowitz MVO: Classical mean-variance optimization with covariance estimated over a 252-day rolling window.
- Risk Parity: Equal marginal risk contribution via convex optimization.
- Equal Weight: Uniform 10% allocation; rebalanced monthly.
- Buy & Hold (Index): Representative of passive S&P 500 tracking.
- DDPG: Deep Deterministic Policy Gradient with continuous action space.
- PPO: Proximal Policy Optimization using the same state space as DQN-PO.

5.3 Evaluation Metrics

The performance is evaluated based on: Sharpe Ratio (annualized excess return/volatility, $r_f = 1.5\%$), Maximum Drawdown (MDD), Calmar Ratio (annualized return/MDD), Annualized Volatility, and Annualized Return. Statistical significance of difference between Sharpe Ratios is tested using block bootstrap approach (Politis & Romano, 1994) for $B = 10,000$ samples [14].

6. Results and Analysis

6.1 Overall Performance

Table 4 provides a detailed out-of-sample evaluation for all considered methods during 2020-2023 test period. The method DQN-PO demonstrates the highest Sharpe ratio (2.14), annualized return (38.6%), and minimum maximum drawdown (-8.3%). Calmar ratio is 4.65, which is significantly higher than the second-best method DDPG (Calmar 2.80).

Table 4: Out-of-Sample Performance Comparison (2020–2023) — All Methods

Algorithm	Sharpe Ratio	Max Drawdown	Avg Return/Yr	Volatility	Calmar Ratio	Convergence
DQN-PO (Proposed)	2.14	-8.3%	+38.6%	12.4%	4.65	~800 ep
DDPG Baseline	1.87	-11.2%	+31.4%	13.8%	2.80	~1200 ep
PPO Agent	1.72	-13.5%	+28.9%	15.1%	2.14	~950 ep
Markowitz MVO	1.43	-14.2%	+27.4%	14.6%	1.93	Convex opt
Risk Parity	1.31	-12.8%	+24.7%	11.9%	1.93	Convex opt
Equal Weight	1.18	-19.7%	+28.1%	17.2%	1.63	N/A
Buy & Hold (Index)	1.22	-20.0%	+28.7%	18.1%	1.58	N/A

All statistics based on out-of-sample test period 01/01/2020-31/12/2023. Bold row = DQN-PO (proposed). Risk free rate is used to calculate Sharpe ratio = 1.5% per year. Max Drawdown = peak to trough net asset value (NAV) drawdown. Calmar Ratio = Return Annualized / |MDD|. ♦ $p < 0.01$ versus Markowitz Mean Variance Optimal allocation by block bootstrap method.

6.2 Quarterly Performance Attribution

Table 5 breaks down the performance results according to the calendar quarters of the years 2020 and 2021, taking into account the COVID-19 crash (Q1 2020) and recovery. The strongest relative performance of DQN-PO occurs in Q1 2020 (-8.3% vs -19.7% for Equal Weight), showcasing its adaptive drawdown protection. During the V-shaped recovery (Q2 2020), DQN-PO captured +28.4%, outperforming all baselines.

Table 5: Quarterly Returns Comparison — DQN-PO vs Baselines (2020–2021)

Quarter	DQN-PO Return (%)	Markowitz (%)	Equal Weight (%)	S&P 500 Bench (%)
Q1 2020	-8.3	-14.2	-19.7	-20.0
Q2 2020	+28.4	+18.6	+17.2	+20.5
Q3 2020	+9.1	+6.4	+6.8	+8.5
Q4 2020	+14.7	+10.2	+11.3	+12.2
Q1 2021	+7.8	+4.5	+5.9	+5.8
Q2 2021	+11.2	+8.3	+7.4	+8.5
Q3 2021	+5.4	+3.9	+4.1	+4.7
Q4 2021	+12.1	+9.8	+8.7	+10.9
FULL YEAR 2021	+38.6	+27.4	+28.1	+28.7

Net returns per quarter for DQN-PO and three baselines. Green indicates positive quarters, red indicates negative quarters. COVID crash window is indicated by a bracket (Q1-Q2 2020). The last row contains yearly totals.

6.3 Ablation Study

To measure the contribution of individual architectural components, an ablation analysis is performed by eliminating one component at a time from the complete architecture of DQN-PO. Results are presented in table 6.

Table 6: Ablation Study — Impact of Architectural Components on Sharpe Ratio

Configuration	Sharpe Ratio	Max Drawdown	Avg Return	vs Full Model	Sig.
DQN-PO Full (Proposed)	2.14	-8.3%	+38.6%	—	—
w/o Dueling Architecture	1.93	-10.1%	+34.2%	-0.21 (-10%)	◆◆
w/o Prioritized Replay	1.98	-9.7%	+35.8%	-0.16 (-7%)	◆◆
w/o Double Q-Learning	1.88	-11.4%	+33.1%	-0.26 (-12%)	◆◆◆
w/o Transaction Cost Penalty	1.71	-13.8%	+41.2%	-0.43 (-20%)	◆◆◆
w/o Macro State Features	1.74	-12.9%	+33.0%	-0.40 (-19%)	◆◆◆
Vanilla DQN (No Improvements)	1.45	-16.7%	+27.3%	-0.69 (-32%)	◆◆◆

Ablation study removing one component at a time. ◆ $p < 0.05$, ◆◆ $p < 0.01$, ◆◆◆ $p < 0.001$ by block bootstrap vs full model. Elimination of transaction cost penalty improves raw return, although it has an extremely negative effect on risk metrics.

The two most influential aspects are the multi-factor macro state attributes (-19% for Sharpe ratio upon their elimination), and the elimination of double Q-learning algorithm (-12%). Transaction cost penalty is crucial for implementation as its elimination not only boosts raw returns but also causes huge drawdowns (-20% for Sharpe).

7. Discussion

7.1 Practical Deployment Considerations

The DQN-PO framework was built for institutional use cases involving corporate treasuries and asset management. Inference time at inference stage is less than 2 milliseconds on an NVIDIA A100 GPU, fitting well within the rebalance window of one day. The discretized action space creates a natural audit trail; each portfolio action corresponds to a pre-computed weight vector, thus fulfilling the stipulations of MiFID II Article 25 and SEC Rule 15c3-5. The transaction cost model is conservative, based on the 10-basis point rate per unit, which is the prime brokerage rate at institutional levels for sector ETFs.

7.2 Limitations and Risks

There are several aspects which require clarity in regard to them. For instance, there is the issue of using the universe of only 10 assets when generalization may involve the use of the 500-asset universe, hence the need for action space compression through hierarchical reinforcement learning or continuous action space [17]. In addition, the assumption that the environment under DQN-PO is stationary over the 3 years' period of training is unlikely to hold when considering monetary policy changes that take place after 2022.

7.3 Comparison with Continuous-Action Methods

The superior performance of DQN-PO compared to DDPG (Sharpe 1.87) and PPO (Sharpe 1.72) indicates that proper discretization of actions space has no significant detrimental effect on performance for the daily rebalance strategy. The continuous-action models may have their own benefits when operating at a faster frequency (intraday) or with an expanded universe of assets.

8. Conclusion

In this study, have presented an optimized deep reinforcement learning framework named DQN-PO based on a dueling network structure and several other components such as Double Q-Learning, Prioritized Experience Replay, and Multi-Factor State Representation, and used this algorithm on the portfolio optimization problem by selecting ten US sector ETFs over a tough four-year out-of-sample period, including the impact of the COVID-19 pandemic.

The experimental results show that proposed method outperforms all seven benchmarks when considering risk-adjusted returns as a criterion. Among others, the Sharpe Ratio is 2.14, Maximum Drawdown is -8.3%, and Annualized Return is 38.6%.

Based on ablation experiments, it is confirmed that each element of network structure significantly affects its performance. In particular, Multi-Factor State Space representation, especially the economic factors of VIX Term Structure and Yield Curve Dynamics, provides the highest marginal value compared to other factors.

For future works, aim to develop a Hierarchical Reinforcement Learning method for large-scale asset universes, a Transformer-based State Space Representation to incorporate long-range market dependency, and Meta Learning approaches for quick adjustment to different market environments. Furthermore, intend to conduct interpretability studies, SHAP, and Attention visualization techniques, to comply with regulatory requirements.

References

1. Markowitz, H. (1952). Modern portfolio theory. *The Journal of Finance*, 7(1), 77–91.
2. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533. <https://doi.org/10.1038/nature14236>
3. Black, F., & Litterman, R. (1990). Asset allocation: Combining investor views with market equilibrium. *Goldman Sachs Fixed Income Research*, 115(1), 7–18.
4. Aripin, N., Hussin, M. R. A., Amran, N. A., Ahmad, H. N., & Ahmad, M. S. (2025). Beyond profit: The role of CSR in enhancing corporate reputation. *Indian Journal of Information Sources and Services*, 15(1), 188–195. <https://doi.org/10.51983/ijiss-2025.IJISS.15.1.24>
5. Yang, H., Liu, X. Y., Zhong, S., & Walid, A. (2020). Deep reinforcement learning for automated stock trading: An ensemble strategy. In *Proceedings of the First ACM International Conference on AI in Finance* (pp. 1–8). ACM. <https://doi.org/10.1145/3383455.3422540>
6. Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1), 3–56. [https://doi.org/10.1016/0304-405X\(93\)90023-5](https://doi.org/10.1016/0304-405X(93)90023-5)
7. Qian, E. (2005). Risk parity portfolios: Efficient portfolios through true diversification. *PanAgora Asset Management Research Paper*, 1(1), 1–10.

8. Deng, Y., Bao, F., Kong, Y., Ren, Z., & Dai, Q. (2016). Deep direct reinforcement learning for financial signal representation and trading. *IEEE Transactions on Neural Networks and Learning Systems*, 28(3), 653–664. <https://doi.org/10.1109/TNNLS.2016.2522401>
9. Carta, S., Ferreira, A., Podda, A. S., Recupero, D. R., & Sanna, A. (2021). Multi-DQN: An ensemble of deep Q-learning agents for stock market forecasting. *Expert Systems with Applications*, 164, 113820. <https://doi.org/10.1016/j.eswa.2020.113820>
10. Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, 1(2), 223–236. <https://doi.org/10.1080/713665670>
11. Huang, J. (2024). Impact of non-performing corporate assets on shareholder's equity and return on the application of AI and blockchain technologies. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications*, 15(3), 412–423. <https://doi.org/10.58346/JOWUA.2024.I3.027>
12. Moody, J., Wu, L., Liao, Y., & Saffell, M. (1998). Performance functions and reinforcement learning for trading systems and portfolios. *Journal of Forecasting*, 17(5–6), 441–470. [https://doi.org/10.1002/\(SICI\)1099-131X\(199809\)17:5/6<441::AID-FOR704>3.0.CO;2-K](https://doi.org/10.1002/(SICI)1099-131X(199809)17:5/6<441::AID-FOR704>3.0.CO;2-K)
13. Kaur, K., & Chandra, G. (2024). Demographic Data Gaps and the Challenges of Population Modeling in Low-resource Settings. *Progression Journal of Human Demography and Anthropology*, 2(1), 13-16.
14. Politis, D. N., & Romano, J. P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, 89(428), 1303–1313. <https://doi.org/10.1080/01621459.1994.10476870>
15. Van Hasselt, H., Guez, A., & Silver, D. (2016). Deep reinforcement learning with double Q-learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1), 2094–2100. <https://doi.org/10.1609/aaai.v30i1.10295>
16. Ding, Q., Wu, S., Sun, H., Guo, J., & Guo, J. (2020). Hierarchical multi-scale Gaussian transformer for stock movement prediction. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence* (pp. 4640–4646). <https://doi.org/10.24963/ijcai.2020/640>
17. Ivanov, A., Petrova, O., & Pavlov, D. (2025). Quality management data-driven decisions fail and how to fix it. *National Journal of Quality, Innovation, and Business Excellence*, 2(1), 1–10.
18. Faramarzi, M., & Ashrafi, H. (2016). Effect of restatement of financial information on the growth rate of financing in companies listed in Tehran Stock Exchange. *International Academic Journal of Economics*, 3(2), 50–63.
19. Sharpe, W. F. (1966). Mutual fund performance. *The Journal of Business*, 39(1), 119–138. <https://doi.org/10.1086/294846>
20. Fournier, A. (2023). A study on nanomaterials incorporated smart materials for flexible and wearable sensors. *International Academic Journal of Science and Engineering*, 10(4), 20–25. <https://doi.org/10.71086/IAJSE/V10I4/IAJSE1035s>
21. Carhart, M. M. (1997). On persistence in mutual fund performance. *The Journal of Finance*, 52(1), 57–82. <https://doi.org/10.1111/j.1540-6261.1997.tb03808.x>
22. Rajan, C., & P. Dineshkumar. (2025). Embedded Machine Learning Framework for Real-Time Prediction of User Information Needs in Intelligent Systems. *Journal of Integrated VLSI, Embedded and Computing Technologies*, 2(2), 73-79.
23. Leila Ismail and M. Ahmad, "Multi-Objective Evolutionary Algorithms for AI-Accelerated Sub-5 nm Floorplanning", *Electronics Communications, and Computing Summit*, vol. 2, no. 4, pp. 1–11, Dec. 2024.
24. Mil Castiñeira, K. Francis. (2025). Model-Driven Design Approaches for Embedded Systems Development: A Case Study. *SCCTS Journal of Embedded Systems Design and Applications*, 2(2), 30-38.