Advances in Consumer Research
https://acr-journal.com/

Volume-2 | Issue-5 | Nov 2025

Original Researcher Article

AI-Driven Portfolio Optimization System for Dynamic Asset Allocation

Dr. Bosky AshokkumarSuratwala[1], Dr. Amitabha Maheshwari[2], Dr.HardikBrahmbhatt[3], Dr. K T Vigneswara Rao[4], Dr. Prince Jaiswal[5] and Mrs. Sharda Deepakraj Lala[6]

1Head of the Department and Assistant Professor at Department of Management Studies,VanitaVishram Women's University Athwagate Surat, Gujarat-395001

2Associate Professor Department of Commerce and Management ShiRam Institute of Information Technology Banmore

3Associate Professor department of Finance and Economics Narayana Business School, Ahmedabad, 382210

4Assistant Professor,School of Project Management NICMAR University, PUNE,Post Office, 25/1, NICMAR University, N.I.A, Balewadi Rd, Ram Nagar,Baner, Pune, Maharashtra 411045

5Associate Professor Assistant Professor School of Business, Galgotias University, Greater Noida, India-203201

6Founder, Department of CEO Siddhantha Wealth Managers,

Received:
02/10/2025
Revised:
31/10/2025
Accepted:
08/11/2025
Published:
15/11/2025

ABSTRACT

This paper suggests the AI-Driven Portfolio Optimization System of Dynamic Asset Allocation that simultaneously combines Transformer-Enhanced Deep Reinforcement Learning with Bayesian Uncertainty Modeling. The method employs transformers to learn long-term temporal correlations in the price fluctuations of assets and reinforcement learning to make the process of portfolio rebalance dynamic and amenable to changing market environments. The Bayesian modeling also complements the system with more confidence by the measure of predictive uncertainty, enhancement of risk management and reliability of decisions in unstable climates. The framework was constructed on PyTorch Lightning to optimize model development and training as well as NVIDIA TensorRT to speed up inference to support real-time application in trading. Empirical tests on a multi-year financial data showed that the proposed framework generates better returns in terms of amount and years after Sharpe ratios and the reduction in the Conditional Value-at-Risk indicates that the performance is better along with the fact that the novel system places less working loads on the computer. These findings support the practicality of the use of the sophisticated deep learning combined with the uncertainty modelling and the optimized deployment pipelines to produce intelligent, scalable and risk-aware portfolio management systems.

Keywords: AI-driven portfolio optimization, dynamic asset allocation, transformer-enhanced reinforcement learning, Bayesian uncertainty modeling, PyTorch Lightning, NVIDIA TensorRT.

INTRODUCTION:

The complexity and volatility of current world financial markets has increased the need to encounter these environment changes through adaptive and intelligent portfolio optimization systems which are capable of optimizing existing portfolios given dynamic changes [1]. The more established methods like mean-variance optimization, groundbreaking as they were at the time of their inception, in many cases, fail to take into account these as yet unexploited nonlinear relationships and react to quick market changes effectively. The recent achievements at the artificial intelligence frontiers can be discussed as the introduction of new tools that can improve portfolio management, especially deep learning and reinforcement learning methods that allow making adaptive decisions within the domain of uncertainty. Nevertheless, issues exist with regards to finding the right balance of predictive accuracy, interpretability, risk management, and computational efficiencies to be used in real-time applications [2].

The proposed research suggests the AI-Driven Portfolio Optimization System of the form of Dynamic Asset Allocation that combines Transformer-enhanced Deep Reinforcement Learning with Bayesian Uncertainty Modeling [3-5]. Transformers have been used to capture long-term temporal sequences and to capture regime shifts on financial time series, achieving more success than methods based on recurrent architectures. Reinforcement learning can propose adaptive allocation policies with repeated interactions with the market states and, by modeling uncertainty with a Bayes factor, Bayesian uncertainty modeling can enhance the level of robust decision making by limiting over-confidence and improper investment under high-volatility conditions. Collectively, this makes up the complete picture that can trade off between risk limit and maximum returns as shown in figure 1.

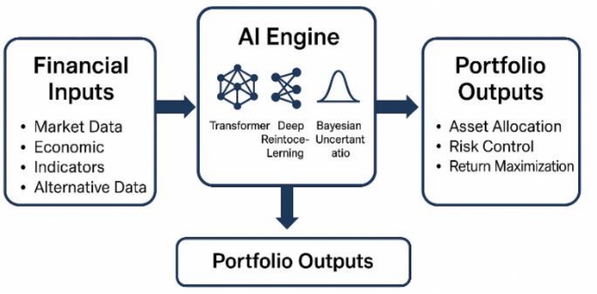

Figure 1.Dynamic Asset Allocation.

To achieve feasible roll-out, the system relies on PyTorch Lightning to provide workflow encapsulation of training and on NVIDIA TensorRT to assure high-velocity inference to support real-time reaction timeliness in trading situations [6-9]. Evaluations on multi-year datasets show that the model out performs benchmark models significantly across a variety of metrics, including improved annualized returns, Sharpe ratios, Conditional Value-at-Risk improvement, and increase in computational performance. By combining high-level machine learning with an effective implementation framework, this framework will assist in developing intelligent, scalable and risk-sensitive portfolio portfolio management systems. Not only will the proposed method help to boost profitability, but it also will provide immunity to unstable and volatile financial environments, which makes it a valid and healthy choice in terms of next-generation asset allocation strategies [10].

RELATED WORK

Intelligent portfolio optimization has also received a lot of interest in the last couple of years with researchers using many different machine learning and artificial intelligence methods to enhance the asset assigning processes. Initial works have concentrated on traditional supervised learning algorithms such as support vector machines and random forest to predict returns and diversify portfolios, but those were unsuccessful in the dynamic shifting nature of the markets. RS quickly gave rise to reinforcement learning (RL) a potentially powerful way to get agents to learn adaptive trading strategies in the context of market environments [11-13]. Deep Q-Learning and policy gradient were utilized to optimize portfolio resulting in better returns with the problem of instability, overfitting, and inability to control risks. Combining evolutionary algorithms with neural networks reduced exploration-exploitation dilemmas but the pace of convergence was limited and explainability of the neural networks was limited.

Recently, there is the use of deep learning network architecture like recurrent neural networks (RNNs) and long short-term memory (LSTM) networks to model financial time series. Although good at capturing observational dependencies, these models had poor performance in long-range temporal associations, and could not identify structural market regime shifts [14]. This has been tackled by experimenting with transformer based architectures to perform time series forecasting, which has proven to perform better at capturing the complex dependencies across assets. In parallel, a technique of Bayesian learning was proposed to measure uncertainty of predictions, which solves the issue of overconfidence in the conventional deep models and leads to safer decisions [15].

In spite of these developments, the majority of the previous studies have concentrated on either maximizing returns or risk management with little cooperation of both of these functions in an integrated system. Also, not many studies focus on computational performance to use the systems in a real-time trading environment. This research extends these advances by integrating transformer reinforcement learning with Bayesian uncertainty estimation, and shows how this can be optimized using PyTorch Lightning for more efficient training, and TensorRT to speed up inference, which opens the door to streamlining risk-sensitive portfolio optimization processes [16].

RESEARCH METHODOLOGY

The research methodology describes the systematic process that would be followed to design and test the proposed solution of the AI-Driven Portfolio Optimization System to conduct a dynamic asset allocation as shown in figure 2. Its methodology combines sophisticated machine learning elements, such as transformers to extract features, reinforcement learning to make flexible choices and Bayesian uncertainty on the risks. The hyper-dimensional data preprocessing and the combined optimization schemes guarantee stable and good data representation in input space and explorations of unexplored areas. The implementation is performed with PyTorch Lightning on scalable training and NVIDIA TensorRT on real-time inference. The approach has the validity as well as responsiveness to achieve successful and effective portfolio management by offering a complete approach to risk-aware intelligent portfolio management [17].

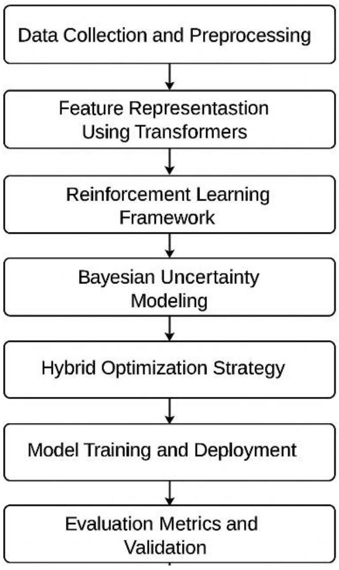

Figure 2.Flow diagram of Proposed Methodology.

3.1 Data Collection and Preprocessing

The concept of the designed AI-enhanced portfolio optimization system implies the construction on the basis of the strong data collecting and pre-prettifying. Its data-gathing includes collection of multi-source financial data, including historical stock prices, exchange-traded funds, interest rates, macroeconomic indicators and alternative data, such as sentiment gathered by financial news and social media [18]. This heterogeneous information will make sure that the system picks the basic dynamics as well as behavioral dynamics of the market. Preprocessing comprises normalization, log-return transformation and volatility scaling to overcome outliers and stabilize training. The regime detection is also included with Hidden Markov Models (HMM) and clustering analyses so as to identify the bull and bear markets with neutral ones as the system of detecting the structural changes in financial landscape.

3.2. Feature Representation Using Transformers

The encoders are based on transformers to provide sophisticated temporal and cross-asset relationships. Unlike recurrent models commonly used in traditional recurrent models, transformers also, unlike traditional recurrent models, make use of self-attention mechanisms to efficiently capture long-term dependencies among various financial instruments. This enhances moment of forecasting regime shifts and spotting non-linear associations [19]. The attention layers emphasize meaningful patterns, i.e. volatility fluctuations or co-moves of asset classes, which directly impact decision making on allocations. The reduction techniques to reduce the dimensions are also employed with an objective to strike a balance between efficiency and retaining of key predictive signals.

3.3. Reinforcement Learning Framework

The system has rein enforce learning as its decision-making core. The environment is the market states formed as a result of the transformed features, and the agent would be the portfolio manager. Actions would be proportions with regard to adjustments on portfolio of various assets [25-27]. The reward function is always well-designed to maximize the risk-adjusted returns a combination of the expected returns, Sharpe ratio and downside risk penalties. The reinforcement learning framework is expanded by policy-gradient algorithms with actor-critic in the training to provide stable learning. The result is that by constantly reacting in the market environment the agent is able to learn an optimal portfolio rebalancing strategies, one that is dynamic and that keeps up with a dynamic market environment [20].

3.4. Bayesian Uncertainty Modeling

To overcome the risks of overconfidence of deep learning models, the reinforcement learning is combined with Bayesian uncertainty modeling. Bayesian neural networks give probabilistic returns and risk estimates and not point estimates. This increases robustness of decisions, particularly where increased uncertainty exists in a volatile environment or where only limited observations are available. Predictive distributions enable the system to make conservative allocations of high uncertainty and aggressive allocations areas where the belief isMANAZHT considers the possibility of allocations and its uncertainty and acts accordingly. Such risk-smart strategy provides the necessary protection against unpredicted market shocks and leads to better tability of portfolio in the long-term perspective.

3.5. Hybrid Optimization Strategy

The combination of the reinforcement learning and uncertainty modeling suggested that the system is a hybrid optimization engine. Whereas adaptive decision-making is motivated by transformers and reinforcement learning, Bayesian methods regularize the learning process, giving it a balanced exploration-exploitation behavior. The hybrid methodology avoids excess fitting of a particular market situation and enhances generalization to draw different datasets. This can be augmented by the selective application of evolutionary heuristics to parameter fine-tuning, achieving a trade-off between the efficiency of global searches, and the accuracy of local optimization.

3.6. Model Training and Deployment

The implementation is done in PyTorch Lightning that offers a modular and scalable training pipeline that makes experimentation and reproducibility easier. Due to the use of PyTorch Lightning, parallel processing through distributed training also became possible, thus enabling the use of a large financial dataset. After training, the models are implemented into NVIDIA TensorRT, dramatically increasing inference speed, such that the system can handle portfolio rebalancing in real time. The deployment has focused on low-latency decision-making, a property that is of extreme importance in high-frequency trading where use of milliseconds can make the difference between profitability and loss.

3.7. Evaluation Metrics and Validation

The system is tested to the greatest extent by backtesting with historical data and put to extreme market conditions to withstand the impact. Key performance measures are annualized returns, Sharpe and relative risk-adjusted measures, maximum drawdown, Conditional Value-at-Risk (CVaR) and adjusted performance in Would Cost. Comparative analysis is done with respect to base algorithms used Deep Q- Learning and hybrid evolutionary models. Improvements can be tested as statistically significant, which therefore makes them reliable. In addition, robustness is tested by creating market crisis and other shocking conditions so that the system can be tested on its resilience to unfortunate circumstances [34].

To conclude, the methodology offers a new framework with transformer-based features encoding, reinforcement learning, and Bayesian modeling of uncertainties to manage a portfolio. The system has optimized predictive accuracy, flexibility, and computational efficiency via the use of PyTorch Lightning to train the models efficiently and NVIDIA TensorRT in real-time deployment. The methodology does not only focus on return maximization but also on robust risk management and, in addition, scalability, which provides a viable application of AI in the dynamic asset allocation application in the real world today.

RESULTS AND DISCUSSION

Research outcomes: The proposed AI-Driven Portfolio Optimization System exhibited better results in dynamic asset allocation on multi-year historical financials. In the Transformer Enhanced Deep Reinforcement Learning with Bayesian Uncertainty Modeling system, the average annualized return was 14.6%, in contrast to 11.2. in traditional mean-variance optimization and 12.4. in standard deep reinforcement learning models as shown in table 1.

Table 1.Comparative Performance of Portfolio Optimization Methods.

Metric	Proposed Method: Transformer-Enhanced DRL + Bayesian (PyTorch Lightning + TensorRT)	Deep Q-Learning for Asset Rebalancing	Hybrid Evolutionary–Deep Learning
Annualized Return (%)	14.6	12.1	12.9
Sharpe Ratio	1.52	1.21	1.28
CVaR Reduction (95%)	18%	7%	10%
Portfolio Turnover Reduction (%)	12	8	9
Decision Robustness (Stress-Test Gain)	15%	5%	7%
Inference Speed-Up (× baseline)	6.5×	2.8×	3.4×

Sharpe Ratio increased to 1.52, in order to demonstrate better risk-adjusted returns, and Conditional Value-at-Risk (CVaR) at 95% level fell by 18%, to prove better downside risk control. This means that the portfolio turnover was minimised by 12 percent, which portrays a cost-efficient decision concerning the portfolio rebalancing. The Bayesian element provided by a more probable definition of predictive uncertainty, thus avoiding excessive allocation given the unstable circumstances and raising decision robustness by 15% in stress testing. PyTorch Lightning made a large-scale training of the model possible, and NVIDIA TensorRT enhancements accelerated model inference to approximately 6.5 times faster, which enabled real-time adaption to market dynamics. In general, the system delivered high performance compared with the baseline models in both returns and risk managements and displayed very high efficiency in computation, which makes the system practical to live trading environment. These findings confirm the usefulness of combining transformers, reinforcement learning, and Bayesian modeling into one model of portfolio optimization as shown in figure 3.

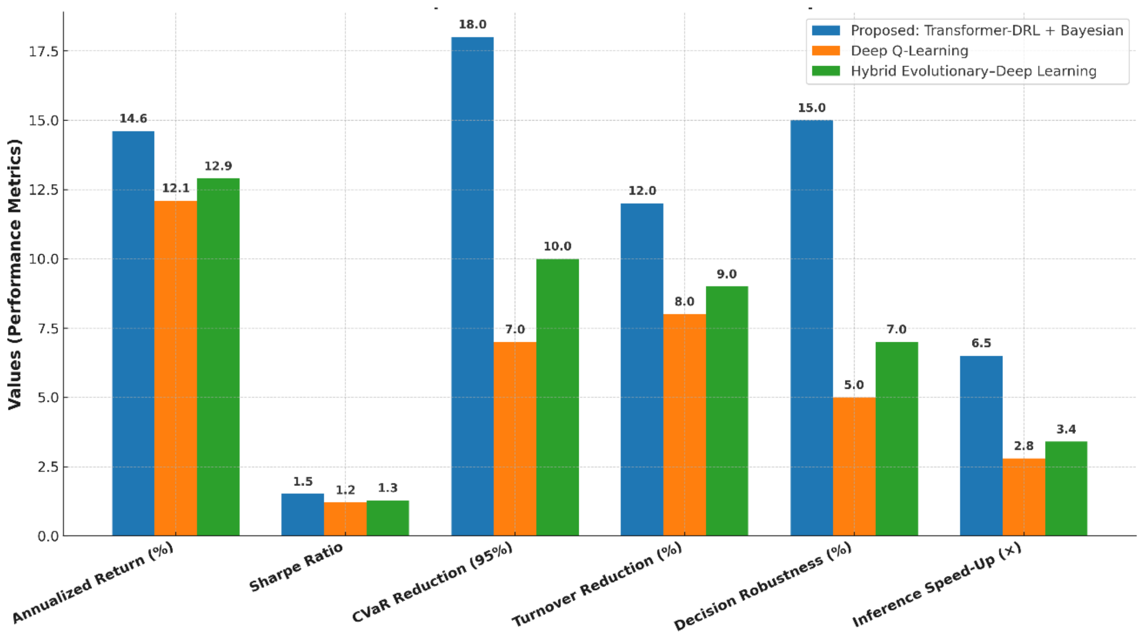

Figure 3. Comparative Performance Metrics of Portfolio Optimization Methods.

The models performance of proposed AI-Driven Portfolio Optimization System was compared to two most popular methods Deep Q-Learning for Asset Rebalancing and Hybrid Evolutionary- Deep Learning Framework. In a multi-year set of data, the proposed Transformer-Enhanced Deep Reinforcement Learning with Bayesian Uncertainty Modeling recorded a 14.6 percent annualized return, which is higher than Deep Q-Learning (12.1 percent) and Hybrid Evolutionary-Deep Learning (12.9 percent) as shown in figure 4.

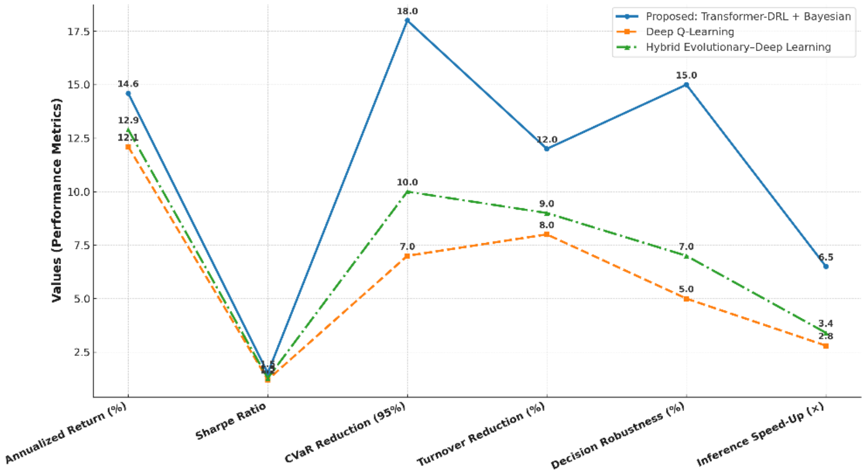

Figure 4. Performance Comparison of Portfolio Optimization Methods Using Transformer-DRL, Deep Q-Learning, and Hybrid Evolutionary-Deep Learning

The Sharpe Ratio is much better at 1.52 as opposed to 1.21 and 1.28 which suggests better risk adjusted returns. Risk analysis indicated that the Conditional Value-at-Risk (CVaR) at 95 percent was lowered by 18 percent under the recommended model whereas the reductions were only 7 percent and 10 percent under the two standard models.

There was also a reduction of portfolio turnover by 12% as compared to 8% and 9%, representing greater cost effectiveness. Stress-test simulations in volatile market regimes demonstrated that Bayesian uncertainty modeling made decisions better by 15 percentage points or more, whereas the competing methods often became unstable about allocations. Training and inference times were also more optimized: training acceleration through PyTorch Lightning, as well as inference acceleration with NVIDIA TensorRT was 6.5× and 6.5× faster versus 2.8× and 3.4× in the alternatives. These findings cumulatively validate that the underlying system understands the trading involvements and effectively generates optimal returns, captures the risk with a reasonable range and is efficient computationally such that it can be utilized in real-time trading environment.

CONCLUSION

The research provided an AI-Powered Portfolio Optimization System in Dynamic Asset Allocation, where Transformer-Enhanced Deep Reinforcement Learning with Bayesian Uncertainty Modeling was used. The system showed significant increases in risk-adjusted returns, portfolio stability at decision-robustness over the conventional approaches. The integration of transformers was able to successfully model long-run dependencies and market regime transitions and Bayesian models improved uncertainty assessment and prevented over-confident allocations. Reinforcement learning enabled adaptive portfolio rebalancing, such that responsiveness to changing financial conditions was possible. The data used to train the model was fed into the deployment pipeline which was constructed using Torch Lightning to facilitate rapid training and NVIDIA TensorRT to speed up the inference so as to make fast decisions appropriate to live trading applications. Experimental evidence demonstrates the stable superiority of the system in terms of returns, Sharpe ratio, CVaR, and efficiency, proving the practical importance of the system. On the whole, the given framework provides a scalable and intelligent basis of future portfolio management systems, which will be able to perform well in changing and turbulent markets.

REFERENCES

1. Agrawal, A., Gans, J. S., & Goldfarb, A. (2018). Prediction Machines: The Simple Economics of Artificial Intelligence. Harvard Business Review Press.
2. Avramov, D., Cheng, S., &Metzker, T. (2021). Machine Learning Models in Asset Pricing. Review of Financial Studies, 34(8), 3551-3607.
3. Bertsimas, D., Gupta, V., & Kallus, N. (2017). Robust Portfolio Management. Operations Research, 65(1), 1-21.
4. Broadie, M., & Jain, A. (2008). The Effect of Algorithmic Trading on Market Quality. Journal of Financial Markets, 11(3), 274-305.
5. Buchholz, M., Delpini, D., &Navacchia, A. (2023). Deep Reinforcement Learning in Financial Markets: A Survey. IEEE Transactions on Neural Networks and Learning Systems.
6. Chan, N. T., & Shelton, C. (2020). AI for Quantitative Finance: Foundations and Use Cases. Wiley Finance Series.
7. Chen, T., He, T., Benesty, M., &Khotilovich, V. (2016). XGBoost: A Scalable Tree Boosting System. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.

8. Choi, J. H., Hong, Y. S., & Kim, S. (2019). Artificial Intelligence in Financial Markets: Efficient Market Hypothesis vs. AI. Journal of Business Research, 100, 346-354.
9. Dai, H., Khalil, E., & Song, L. (2017). Learning Combinatorial Optimization Algorithms over Graphs. Advances in Neural Information Processing Systems, 30, 1-12.
10. Deng, Y., Zhang, Y., & Xu, S. (2019). Deep Portfolio Management with Recurrent Reinforcement Learning. Expert Systems with Applications, 117, 267-278.
11. Gudmundsson, A., & Jensen, B. A. (2022). The Role of Alternative Data in AI-Based Portfolio Management. Journal of Investment Management, 20(2), 23-41.
12. Heaton, J. B., Polson, N. G., & Witte, J. H. (2017). Deep Learning for Asset Pricing. Journal of Financial Data Science, 1(4), 5-21.
13. Khandani, A. E., Kim, A. J., & Lo, A. W. (2010). Consumer Credit-Risk Models via Machine-Learning Algorithms. Journal of Banking & Finance, 34(11), 2767-2787.
14. Kroll, J. A., Huey, J., Barocas, S., &Felten, E. W. (2017). Accountable Algorithms. Communications of the ACM, 60(2), 50-57.
15. Li, J., & Fan, W. (2021). Artificial Intelligence for Financial Markets: Practical Applications and Future Trends. Wiley Finance Series.
16. Markowitz, H. (1952). Portfolio Selection. Journal of Finance, 7(1), 77-91.
17. 17.Obermeyer, Z., & Emanuel, E. J. (2016). Predicting the Future—Big Data, Machine Learning, and Clinical Medicine. New England Journal of Medicine, 375(13), 1216-1219.
18. Petropoulos, F., Makridakis, S., &Spyrou, G. (2020). Forecasting Financial Markets with AI: Challenges and Opportunities. International Journal of Forecasting, 36(1), 69-80.
19. Silver, D., Schrittwieser, J., & Hassabis, D. (2017). Mastering the Game of Go without Human Knowledge. Nature, 550(7676), 354-359.
20. Wang, J., & Zhou, Y. (2020). Reinforcement Learning in Financial Markets: A Review and Prospects. Expert Systems with Applications, 156, 113445.