

# **Reinforcement Learning and Stochastic Optimization in Multi-Asset Portfolio Construction: A Data-Driven Investment Framework**

*Dr. Marko Bohanec, Associate Professor of Computer Science, University of Ljubljana, Slovenia*

---

## **1. Introduction**

Artificial intelligence (AI) today is critical in portfolio optimization. Since the 1960s, technology has become a significant part of investing. The way investment managers generate outperformance has completely changed. Advancements in big data, machine learning, and AI have allowed managers to uncover innovative ways for the detection of valuable insights that human capability could not do so easily. Portfolio optimization is closely coupled with asset allocation, which is one of the most critical processes for investment management. The more sophisticated the tools and approaches used in portfolio optimization, the more desirable results are expected.

The term "AI-driven approach" is related to how a range of AI technologies could be used to achieve a specific goal or to solve a problem. For instance, in making decisions, an AI-driven approach can be used in prediction, classification, and estimation. This research study aims to identify contemporary AI techniques and tools for use in portfolio optimization and will assess contemporary AI algorithms and the performance of these by measuring the relative performance of two AI algorithms with the traditional portfolio optimization algorithms. This is because the traditional methods have limitations. In a typical portfolio optimization, it is challenging to simultaneously optimize real-time transaction costs and the illiquidity risk for a frictionless problem. It is also difficult to achieve satisfactory results directly in the robust portfolio optimization problem.

This limitation of using traditional methods to solve these portfolio optimization problems opens great potential for AI. Especially, reinforcement learning has gained more attention in recent years because it allows the agent to learn from the environment

and take the actions that maximize the cumulative reward. It is also suitable for a broad range of optimization problems, financial ones in particular. With the aid of reinforcement learning, we have succeeded in effectively optimizing an investment fund, with initial capital allocation in a variety of stocks.

## **1.1. Background and Significance**

### **1.1.1. Context**

The concept of portfolio optimization comprises fundamental research in the evolution of investment strategies. While value, growth, index, and other approaches have highlighted traditional investment, many recent efforts have concentrated on the algorithmic and AI-driven vision. Known as active or dedicated management, this work seeks to create and then operate one or more optimized portfolios, for example, portfolios with high returns, small risks, or both combined. Optimization need not be separated into two connected activities, as can be seen in active management contexts: (i) portfolio construction (portfolio setting) and (ii) portfolio change at the time. In this research, we focus on the first situation, where the portfolio construction (optimization) process is designed to be deterministic, ranking-based, and need not change.

Significant advancements in specialized knowledge appear in investment management, such that AI-driven decision-making systems comprise computational networks that focus on simulating various investment strategies to illustrate the management of fundamental behavior, to offer advice to agents or shareholders on how to take action or engage in investment activities. While some works study various aspects of decision-making systems, especially their effect on market behavior and efficiency, our work looks at computational networks for educational purposes. Such networks, which are of appropriate architectures in current machine learning technologies, have already been widely applied in the field of finance, as in what is commonly known as algorithmic trading. The algorithm, even if it is left to work and perform a particular strategy, has not yet been used to determine financial decision-making in terms of a diversified portfolio. Moreover, it has not been used in AI financial decision-making that specifically seeks to optimize minimization, maximization, or arbitration returns when it is coupled with limiting risk. This will be a critical addition to the finance and artificial intelligence literature, and it is of great premier and topical importance. The machine learning subcategory is the connection network for this topic. This approach has been

used with satisfactory results in various research, and in our work, thematically, the networks can help us understand and accomplish the task.

## **1.2. Research Objectives**

1.2. Research Objectives. The main purpose of this study is to verify the potential of AI-driven techniques in portfolio management, specifically if there are any statistically significant differences in comparison with the traditional model in terms of the portfolios selected by the evaluated classifiers. Additionally, we want to check if using AI-driven techniques yields any statistically significant differences in comparison with one another. Having stated the primary purpose, the following specific objectives can be listed: to determine if there are any statistically significant differences between the portfolios built using the classical mean-variance optimization and those selected by the evaluated classifiers or the preferred alternative similarity measure; to determine if there are any statistically significant differences between the portfolios selected by the evaluated classifiers using the classical mean-variance optimization and those constructed using the evaluated classifiers; to determine if there are any statistically significant differences between the portfolios built using different models and those selected by the evaluated classifiers or the preferred alternative similarity measure; to determine if there are any statistically significant differences between the portfolios selected by the evaluated classifiers using different models.

The reason for examining the above-mentioned objectives is the assumption that by using machine learning-enhanced models (along or instead of more conventional approaches), it is possible to obtain better results in terms of portfolio performance in relation to classical investment strategies. We take as the starting point the assumption that the traditional model, resulting from the classical optimization problem solution, is a good one, but based on technological progress, there is a constant search for even more effective procedures. It seems interesting to compare the performance in terms of different measures for the candidates of the classification algorithms and AI-aided models created based on sensible classifiers, as well as those that are the fruits of hybrid systems, with pure AI-aided ones. The essence of this approach is not pointing out which classifier performs best in constructing the investment position. The point is, however, to determine the improvement rate by a rigorous comparison of the models

constructed against the same nature of the input data. Each of these points is documented by means of a statistical approach.

## **2. Foundations of Portfolio Optimization**

Modern finance has its roots in portfolio optimization, the idea of balancing risk and return to best grow wealth. One of the investors' chronic problems is the uneasy trade-off between the return from assets and their risk, primarily measured by their volatility of returns and their covariance. Traditionally, the goal of portfolio optimization is to find a fraction of an investment's wealth to be placed in each available asset under management. The optimization is based on historic returns of assets and return expectations to minimize loss while maximizing rewards. The historic returns of assets used to forecast returns are usually obtained by analyzing assets from financial data such as stocks. Over the last six decades, a considerable effort was devoted to finding optimal portfolio strategies from various perspectives, in particular through optimal algorithmic methodologies for designing portfolios.

The conventional trade-off between risk and return is to diversify away from firm-specific risk by investing in many stocks. An investor in a well-diversified portfolio is safe from the idiosyncratic risk of any of these companies. Finance theorists often cite portfolio theory. This portfolio theory demonstrates modern portfolio theory and its concepts, like the differences between risk and volatility; that risk is the portfolio standard deviation because it represents both systematic risk and idiosyncratic risk; the role of diversification in reducing both risk and volatility; and that correlations represent part of risk, identified further by using the concept of non-marketable bonds. Modern investment strategies are all generally based on this theory, but for many reasons, managers are hesitant to utilize modern mean-variance optimization in marketing and practical investing. A large yet incomplete literature documents the rejection of frontier investing. At the same time, constraints and transaction costs are considered. Minimum-variance optimization, in contrast to the frontier, remains a popular starting point for portfolio optimization. Simulation has been used to argue that these constraints and transaction costs render mean-variance barely an improvement over optimization. Diversification is a focal point within the theory. Empirically, convergence with mean-variance optimization is examined theoretically. Portfolio theory has helped in implementing diversification optimization in practice even if mean-variance

optimization has not made a significant impression. Implementing modern mean-variance optimization is difficult since its empirical tests are poor, and costs are impactful. The fundamental logic of this theory readily led to mean-variance analysis, which is once again drawn upon for the direct treatment of optimization needed for hedging. Modern portfolio theory and its diversification concepts have certainly influenced the investment community in the practice of the compartmentalization of broad markets. Each portfolio is designed for certain investor types or objectives, all linked by similarities that could be general representations of the established dilution objective. Modern diversification is at the origin of much contemporary investment practice. Today, combinatorial optimization is frequently used when diversifying a portfolio of assets. Modern financial theory is the basis of a wide variety of current investment strategies.

### **2.1. Modern Portfolio Theory**

The fundamentals of investing and portfolio dynamics are analyzed in Modern Portfolio Theory. Among the essential attributes of MPT is the trade-off between risk and return. In essence, MPT posits that higher returns are accompanied by higher levels of risk. This strategy aids market participants in making investment choices that are in line with their risk and return expectations. Moreover, MPT defines the efficient set of portfolios that provide the lowest level of risk for a given return, or the highest possible return for a given level of risk. The underlying mathematical calculations support the existence of these portfolios and the quantitative metrics used to describe and contrast them. While stocks and funds have the greatest potential for higher returns, they are also the most volatile and unsafe investments available.

Diversification, or the spreading out of investments, is a popular approach to mitigating the risk of single-stock investments. Essentially, MPT increases the likelihood of a portfolio's performance moving towards an average weighted result. The reliance on the principles established in MPT has weathered its historical criticism. As AI becomes more advanced and accessible, it is emerging as an alternative technique. But as we progress, it is necessary to address MPT's mettle and relevance concerning the modern investor community.

## **2.2. Efficient Frontier**

Efficient frontier (EF) is an essential concept in investment management, which shows a range of portfolios with the highest returns for a given level of risk. The theory states that this set of portfolios has the maximum expected return for every level of risk or the minimum risk for every level of return. It is the foundational piece of modern portfolio theory. The importance of the EF lies in the fact that it allows investors to quantify the relationship between risk and return. The generated EF chart assists investors in choosing an optimal portfolio, which maximizes the expected return for a given level of risk or minimizes the risk for a given level of return. In classical finance, investor preferences are captured by modern portfolio theory or other forms. However, these traditional forms of investor preferences restrict the commercial flexibility of the reconstituted portfolio, which is disadvantageous. In the portfolio construction stage, investors are unable to choose any portfolio with any level of risk since they can only select portfolios within the depicted region. In the current status, market dynamics have presented rapid fluctuations due to increasing information volume and velocity, thereby exaggerating the findings of traditional forms of EF. For more accurate measurements in different market phases, AI algorithms have been deployed to construct up-to-date breakthroughs in portfolio optimization. The invention of the AI-Portfolio theorem provides the argument that using AI to select the frontier is better than using a frontier consisting of traditionally reconstituted portfolios for future outcomes. AI can redefine EFs to identify new and better frontiers in dynamic intervals.

## **3. Role of Machine Learning in Investment Management**

Just as machine learning revolutionized processes in many other domains, it demonstrates the potential to change investment management. Investors always seek ways to make better predictions or improve their understanding of the market, for instance, by recognizing patterns from historical or real-time data of financial firms. Machine learning algorithms outperform traditional data analysis, particularly in the application field of investment; this is most useful for financial portfolio optimization. Machine learning algorithms can model complex multidimensional data effectively and yield multi-level optimizations. In this domain, portfolio optimization can be viewed as selectively determining the best investment options under conditions of multiple investments, and the selection criterion may reflect trade-offs concerning maximizing expected return and minimizing risk measures.

Many core machine learning algorithms have proven to be very effective in the field of finance. For instance, decision trees have applications in both time-series and non-time-series problems, which can also perform well when multicollinearity is present within the dataset. Random forests utilize several decision trees, and this method offers a certain protection against overfitting and increases output accuracy. Support Vector Machine generates a hyperplane that has the largest distance to the closest data point of any class; this method is good at classification and regression based on statistical learning theory. Artificial Neural Network is used for technology forecasting and risk management, specifically for predicting company daily returns with user-generated stock characteristics and predicting asset returns. Gradient Boosting Decision Tree integrates several decision trees, such that the current decision tree output is regressed by optimizing the relationship between current and the residuals of previous decision roots. However, these algorithms, particularly Artificial Neural Network and Gradient Boosting Decision Tree, are well-known for their good performance in finance but are complex and not easy to interpret. In either time-series or non-time-series applications, the dataset is large with latent factors and noise. These complex factors may be linked closely, and it is hard to reveal a clear relationship. Given these challenges and potentials, there is scope for research and application in the field of finance in quantitative ways using machine learning. The main characteristics of this domain include being non-linear, ubiquitous with noisy and complex data, and large scalability or dimensionality. These characteristics also pose challenges and require advanced data analysis to provide insights into the development of innovative methods and techniques. In the following, we focus on a portfolio management form to structure the survey, given it is a typical investment scenario and touches upon the essential topics of retail investors or financial managers.

### **3.1. Types of Machine Learning Algorithms**

Machine learning algorithms are classified into three types. Supervised learning is widely used in finance problems when a precise mapping function can be established between past datasets and predicted targets. Quite the opposite, unsupervised learning settings are deployed when there is no target to pair entries, such as clustering, pattern recognition, and recommendation. From a financial viewpoint, implementing unsupervised learning benefits from shallow data with pattern definitions, from fraud detection to economic cycles. Reinforcement learning has attracted more attention in

recent times due to the practical study of ML models. Accommodated parameter controls allow for the continuous adjustment of investment portfolios to quickly shifting economic situations, encapsulating relevant functions that enhance the certainty of surpassing the average controlled adjusted investments. Investors consistently present finances to refine settings.

ML algorithms are superior to traditional signal processing methods because of the knowledge gained throughout these approaches. The performance of most ML methods may be improved by increasing the degree of training used from a past dataset during a supervised learning process. Feature selection, model training, and evaluating the algorithm's predictive power are all crucial steps in the supervised learning model. When implementing ML algorithms, however, there are limitations. ML algorithms may be difficult to interpret and use, so their computational cost and system complexity need to be carefully evaluated in order to guarantee that the information gleaned from the model is valuable. The effective transmission of recommendations, accountabilities, and sources of uncertainty to the end investor can also be crucial. In these settings, supervisory, recommender, and combination systems may be used in scenarios such as portfolio optimization.

### **3.2. Applications in Portfolio Optimization**

Machine learning can be valuable in portfolio optimization with a number of applications. In this subsection, we show different use cases for machine learning in investment applications focusing on asset allocation or portfolio selection.

Forecasting. The most common use of machine learning in investment is to build predictive models using training datasets based on historical simulations, with the aim of predicting the expected returns or the expected risks of investments or portfolios in the future. This can help in better asset allocation. An investment strategy based on predictive modeling uses the learning during the back test as well as during the implemented real-time trading.

Dynamic rebalancing. In a rapidly changing world, an effective dynamic portfolio rebalancing strategy can earn extra returns or protection from market downturns. Based on machine learning, one can design dynamic optimization strategies to rebuild the

desired portfolio weights based on the risk factors of portfolios as well as market indicators, news, or sentiment. This will help to build self-regulated portfolios.

If we sum up, machine learning that is currently developed can be quite useful if it is part of a risk factor network and we can build a new asset allocation method or risk management tool. Machine learning can also help us achieve a self-adjustable, dynamic portfolio management method, where we adjust the portfolio based on the market conditions so it can change the investment profile benchmarks. Although asset selection profit is quite elusive, we have found that true signal methods are able to achieve better performance in quantitative asset selection than traditional methods when tested in a simulated real investment environment. We have also found that feature engineering is crucial in asset selection or portfolio optimization with machine learning. We are generally recommended to take a more careful and rule-based approach when applying artificial intelligence to trading and investments, especially when applying it to generate rules on the set of available instruments.

#### **4. Case Studies and Practical Implementations**

Five case studies and practical implementations of AI-driven portfolio optimization

- We provide five different case studies of demonstrated successful applications of AI-driven portfolio optimization across various asset classes, namely US equities, Indian equities, oil prices, and cryptocurrency. All case studies involve institutional asset managers optimizing live capital. The timescales studied vary, ranging from six to thirty-two months. Each case study follows a broad structure: tailoring data preprocessing, strategies considered, strategy implementation, strategy performance, the evaluation of AI effectiveness, and finally a conclusion where lessons learned and best practices are shared. The evidence generated from these case studies suggests that AI modeling shows considerable promise in the development of portfolio management strategies. Evidence shows that implementing AI-driven strategies in a practical context can lead to highly attractive portfolio performance that overcomes the practical barrier of implementation slippage. Altogether, this allows institutional asset managers to offer market-beating risk-adjusted returns via the addition of state-of-the-art portfolios to their product suite.

- All case studies, as with any real-world implementation, detail challenges of integrating machine learning technologies into the workflow of an asset manager. One of the most significant practical implications gleaned from the case studies is the wide selection of strategies that utilize numerous different techniques to harness the broader goal of superior and risk-adjusted portfolio returns. Both in the historical results over a grand average backtest in the region of fifty years, as well as the live management period throughout which the model is exposed to meticulous live trading, results suggest that (after slippage) AI-driven portfolio management leads to a significantly higher proportion of positive outcomes. On average, these outcomes are much higher in terms of return quantities, which span across asset classes and timescales of the case studies. Overall, the expected incidence of investor capital funded by AI-based portfolios typically also exceeds the standard given a considered degree of capital invested for these active management strategies.

#### **4.1. Real-World Examples of AI-Driven Portfolio Optimization**

The applications of artificial intelligence (AI) technology in the capital market and portfolio optimization have received considerable attention. This section first provides some real-world examples of AI-driven portfolio optimization. Then, I will illustrate a real-world study of cross-industry. Portfolio managers could gain more constructive findings when deciding whether to add AI in portfolio management through these studies.

Real-World Examples of AI-Driven Portfolio Optimization: AI drives real-world portfolio construction with machine learning, asset selection, and allocation enhancements, including upper tail risk reduction. Machine learning (ML) is effective at finding investment companies in many sectors, adopting different strategies, asset selection, and portfolio management. This subsection provides some real-world examples of AI-driven portfolio optimization. Many of the effectiveness of their applications is, to a great extent, verified by the relevant results such as returns, risks, and evaluation metrics they achieve. These examples showcase the cross-industry application scenario of respective models involved.

A Case Study of Enhancements to Portfolio Management: Machine learning and AI have received much academic research attention. Beyond academic insights, the industry has also conducted case studies to present the effectiveness of machine learning in the

capital market and portfolio optimization context. Focusing on real-world case studies and applications can provide convincing findings on AI's various applications in finance. The extent to which an industry is a case study, they developed and assessed AI/ML trading strategies by answering several questions about trading strategies, performance, risk management, financial metrics sensitivity, and other related topics. Overall, they provide us with relatively comprehensive cross-industry findings associated with AI/ML trading strategies. With the effectiveness of production AI-improved portfolio management as a goal, we focus specifically on four of them with powerful neural networks. These networks are not only for optimal construction. Instead, we want these optimizations to be effective in action, as well as in various AI designs, optimal objectives, and application contexts. We also highlight the model in its related industry response.

#### **4.2. Challenges and Limitations**

AI-driven portfolio optimization provides a number of promises, yet a number of challenges and limitations have been encountered. First, challenges are driven by the financial market's complexity: data quality is usually not perfect, with outliers, missing values, and noise being common. Moreover, using financial data might result in models' low interpretability due to the issue of spurious correlations or underfitting and overfitting when new data arrives. Additionally, financial markets are non-stationary, changing through time and space. Hence, developing financial and investment strategies might require expert knowledge, as these changes in the market might be driven by neuroeconomic factors. Besides the technical and domain-related challenges, there are peculiar areas of consideration for practice. Regulatory and ethical considerations might also prevent extensive application of AI in practices. Black-box models' regulations require AI-based practices to explain the ingredients and behaviors of AI models, making them less likely to be adopted. Even from practitioners' side, AI models still need a credible approach to evaluate their performance. Problems have been observed where, in excess of business, AI tools did not outperform traditional models. Nonetheless, enabling adaptive dynamic strategies and combining modern AI methods with traditional models might provide a fruitful approach.

## **5. Future Directions and Emerging Trends**

As AI algorithms continue to show potential to outperform humans in routine tasks, investment managers have the possibility to open up new ways of developing investment strategies in response to technological innovations. Despite the rapidly advancing technology, investment managers were using AI to gain insights, improve predictions, and trading algorithms, while also utilizing advanced analytics. Thus, it seems plausible that AI – which is based on robust economic theories, big data, and advanced computational methods – could help provide consistently outperforming portfolios by scanning and decoding immense amounts of data. Despite the limitations and criticisms of various AI methods, AI is still evolving, and the field’s ambitious goals will eventually change investment management.

Certain AI developments could change the alignment of power between investment professionals and clients. As AI’s progressive sophistication has led to it surpassing humans in certain key areas, the implementation of AI algorithms and big data might enhance investment strategies. The trend for big data and analytics can be anticipated to develop in a variety of AI algorithms as it offers valuable and emerging signals that can be used in trading algorithms, either on a stand-alone basis or in conjunction with fundamental data. The predictability of big data might lead to shifts in strategic perspectives within innovative AI investment management companies. This would involve data being valuable not just because a hypothesis is being tested, but because it is being responded to by the market. Finally, AI could lead to investment strategies becoming more transparent as easier development and marketing have led to a greater number of investors gaining access to systematic strategies. Quantitative systems traditionally involved long and slow periods of investment teams testing systems to ensure predictive robustness and/or behavioral bias. If machine learning strategies are used, the portfolio manager can illustrate to clients exactly how a system might respond over a range of outcomes and how it has been responding. Managers and clients are less reliant on complex, long-term back testing, as AI algorithms continue to develop, and this technology will make accessing strategies many times more cost-effective than they are now. Disruptive risk can be a benefit when such model shifts are tied to the market, especially when they are tied to the trend-considering machine. The AI trading algorithm universe has data providers increasingly viewing their data as a platform for end users to build their trading algorithms and systems. This is because for certain

machine learning-based signals, the mathematical code and modeling used by a given researcher to predict can indeed be of value in the marketplace.

## **6. Conclusion**

The aim of this research was to investigate the potential of incorporating intelligence solutions in the process of portfolio optimization. Along the lines of the above speculation, we report that AI technologies are predicted to revolutionize the traditional investment management sector by providing extensive information on securities research, instantaneous trading, information on market timing, integrated solutions, real-time strategies, and other high-level services. However, AI technologies require further scientific interrogation with respect to world financial data to test whether the AI model can be implemented in the above format. As a result, it is necessary to develop, commercialize and update proprietary methodologies that combine both classical and conventional methods with continuous improvement without disrupting the wider field of finance, which requires expertise and high public policy and private equity.

However, some benefits and biases to be aware of include the role of portfolio management and financial performance, potential ethical dilemmas related to using GANs for portfolio wealth development, and the role of financial market integration in affecting potential trading strategies generated by AI models. Therefore, the question remains, will new AI technologies, regardless of the number of contributed benefits, be able to support risk in the financial market in the coming years? If there is any doubt about this concern, some members of the investment market may continue to express their skepticism, providing a useful description of the effectiveness and extensive effects of AI assets in improving portfolio and other asset allocation processes. The discussion calls on researchers in the AI field to cooperate and merge global financial and portfolio data to facilitate the early development of AI schedule recommendations in various global financial markets.