

The Influence of Artificial Intelligence in Portfolio Management and Risk Optimization

Author: Paul Sakr¹

Affiliation: CIRAME Research Center, Business School

Holy Spirit University of Kaslik, P.O Box 446, Jounieh, Lebanon

Abstract:

The current paper discusses how artificial intelligence (AI) has created a paradigm shift in the field of portfolio management and risk optimization and critically evaluates how machine learning, deep learning, reinforcement learning, and evolutionary computing can overcome the structural limitations of classical quantitative models. The traditional construction of a portfolio, using the Markowitz (1952) mean-variance maximization model, is limited by the assumptions of return normalization, covariance stationarity, and good parameter estimates (which are systematically broken in actual markets). Based on groundbreaking empirical research, such as Gu, Kelly, and Xiu (2020) on machine learning to price assets, Fischer and Krauss (2018) on LSTM to predict markets, and Jiang, Xu, and Liang (2017) on deep reinforcement learning to manage portfolios, the paper will synthesize the reported empirical findings to assess the practicality of AI-based. Findings suggest that machine learning models can substantially outperform out-of-sample return predictions as compared to linear factor models; Gu et al. (2020) find that in some specifications, machine learning forecasts can be roughly two times greater than the predictive alpha of regression-based strategies. The authors show that LSTM networks trained on S&P 500 constituents can obtain daily returns of 0.46% in advance of costs and a Sharpe ratio of 5.8, which is much higher than memory-free classifiers (Lutz, 2018). The reinforcement learning agents also have dynamic, cost sensitive rebalancing of portfolios that are more than the fixed benchmark allocations. The paper finds that AI is a true paradigm shift in quantitative investment management, with better risk-adjusted returns and ability to manage tail-risk, even though overfitting, limitations on interpretability, and sensitivity to transaction costs still pose critical issues in practice.

Keywords:

Artificial Intelligence, Portfolio Optimization, Risk Management, Machine Learning, Deep Learning, Long Short-Memory, Reinforcement Learning, Genetic Algorithms, Mean-Variance Optimization, Sharpe Ratio, Conditional Value-at-Risk, Factor Investing, Quantitative Finance.

Introduction: "The What & Why"

Markowitz (1952) introduced the modern portfolio theory followed by Sharpe (1964) with the Capital Asset Pricing Model (CAPM) and Ross (1976) with Arbitrage Pricing Theory (APT) as the theoretical basis of quantitative investment management. Mean-variance optimization characterizes the efficient frontier - the collection of portfolios which offer the maximum expected return, given a certain amount of variance - and offers an objective and rigorous mathematical model of asset allocative decisions. Although this paradigm has great theoretical beauty, the assumptions upon which it is based are systematically broken by the empirical markets. DeMiguel, Garlappi, and Uppal (2009) compared 14 forms of mean-variance optimization on seven empirical studies and found that none of them outperformed a naive equal-weight (1/N) allocation based on out-of-sample Sharpe ratio, an important indictment of classical optimization in practice. Their discussion also shows that to have a reasonable degree of certainty over the estimation of the parameters that optimize a portfolio of only 25 assets, it would take about 3,000 months of past returns, which is impractical to any typical investor. Similar drawbacks are characteristic of risk management tools: Value-at-Risk (VaR), which is widely used under the Basel II regulatory framework, assumes normally distributed returns, and thus underestimates tail risk by a large factor, which occurred during the 2008 Global Financial Crisis (Jorion, 2007).

The contemporary data environment both poses a challenge and an opportunity. Financial markets are now creating large amounts of structured and unstructured data such as high-frequency price updates, earnings call transcripts, satellite imagery, and social media sentiment, which are structurally inaccessible to classical models, which are only designed to use low-dimensional and low-frequency inputs. Machine learning, deep learning, reinforcement learning, and evolutionary computation have become the main analytical answer to this dilemma and are known as artificial intelligence. In contrast to parametric models, AI approaches learn structure directly on-data, non-stationarity, and non-linear interactions, which are not representable by conventional econometrics. The main research question that drives this paper is: how effective are AI-based techniques, based on the published empirical findings, in improving portfolio management and risk optimization as compared to conventional quantitative methods? This paper presents a synthesis and survey of the most impactful empirical experiments that have proven AI to be effective in making returns and building portfolios, analyzes the claimed benefits of certain AI architectures compared to classical models, and critically assesses the real-world constraints and governance issues that limit the use of AI in investment management (Beainy, 2026).

Literature Review: "The story so far "

This paper is a review of the empirical research done by Gu, Kelly, and Xiu (2020), which was published in the Review of Financial Studies. They apply a wide range of models such as Lasso regression, elastic net, gradient-boosted trees, random forests and neural networks with up to five hidden layers to a sample of about 30,000 US stocks between 1957 and 2016 and find that machine learning models have much higher out-of-sample predictability of returns than linear factor models. What they have found is that neural networks and tree-based approaches, in some specifications, can approximately multiply the predictive alpha of the best regression-based methods by approximately one half of the non-linear interactions among the predictors. The prevailing predictive signals observed by all classes of models including momentum, liquidity,

and volatility demonstrate that AI procedures are more effective in extracting signals than spurious trends, and its benefit is as much as well-known financial-economic intuition. Importantly, the analysis deploys stringent out-of-sample methods that involve recursive expanding windows to avoid look-ahead bias, which gives a lot of credence to its findings (Beainy & Kamel, 2025)..

In the European Journal of Operational Research, Fischer and Krauss (2018) implemented Long Short-Memory networks (LSTM) to predict directional changes in price of the entire universe of all S&P 500 constituent stocks between December 1992 and October 2015 - a large, liquid, and survivorship-bias-adjusted universe. Financial time series in particular are well-posed to LSTM networks due to the gating mechanism (including forget, input and output gates) of the network, which addresses the vanishing gradient issue of standard recurrent neural networks, allowing an LSTM model to learn both short-run and long-run temporal dependence. According to Fischer and Krauss (2018), LSTM networks produce pre-transaction-cost daily returns of 0.46% and Sharpe ratio of 5.8, which is much higher than the benchmarks of the classifiers such as random forests (Sharpe ratio about 4.1), standard deep neural networks, and logistic regression (Sharpe ratio about 3.5). It is worth noting that excess returns are also more heavily concentrated before 2010, and that there is some sign of performance deterioration since that time - a result which has significant implications on the sustainability of AI-based alpha and the rate at which liquid markets learn publicly available predictive information in today's competitive environment (Beainy & Kamel, 2023), especially since pricing became more and more dependent on AI (Khalife et al., 2023).

Jiang, Xu, and Liang (2017) present a deep reinforcement learning (RL) framework of the portfolio management process that presents asset allocation as a time-varying repeated decision problem. They have their own architecture, the Ensemble of Identical Independent Evaluators (EIIE) which includes a Portfolio Vector Memory (PVM) mechanism that represents current holdings in the state of the agent allowing dynamic transaction costs to be explicitly modelled. The framework, which learns the history of cryptocurrency prices (Athari et al., 2025) with three variants of the network, namely Convolutional Neural Network (CNN), basic Recurrent Neural Network (RNN), and LSTM is based on an Online Stochastic Batch Learning (OSBL) scheme directly optimizing cumulative logarithmic portfolio return. The RL framework outperforms uniform buy-and-hold, best single-asset buy-and-hold and classical online portfolio strategies, including Passive Aggressive Mean Reversion (PAMR) and the Online Newton Step (ONS), in backtesting. To supplement such deep learning methods, Chang, Meade, Beasley, and Sharaiha (2000) show that genetic algorithms can solve cardinality-constrained portfolio selection NP-hard intractable to classical quadratic programming, as they encode portfolios as chromosomes and encode binary stock-selection vectors as continuous weight vectors, and solve the problem by tournament selection, uniform crossover, and Gaussian mutation. Despite all of this evidence and the development of AI and blockchain technology (Khalife et al., 2024) there are still a number of critical gaps: most of the research assesses pre-cost performance or uses simplified transaction cost models; the interpretability of AI-driven investment decisions has not been sufficiently studied to meet regulatory requirements (Athari et al., 2026); and most research in the field of RL is done within a cryptocurrency or simulated equity setting, which limits its applicability to institutional equity markets (Kamel et al., 2026).

Methodology: "The how we did it "

This research will be based on a systematic review and comparative synthesis research design, which utilizes peer-reviewed empirical research and published simulation outcomes instead of gathering new data. This is the right way to go considering that the paper is six pages long and that empirically validated findings are readily found in the literature. The major studies on which the analysis is based are: Gu et al. (2020) on supervised machine learning in the field of returns prediction; Fischer and Krauss (2018) on the LSTM-based time-series prediction; Jiang et al. (2017) on deep reinforcement learning in portfolios management; DeMiguel et al. (2009) as the classical benchmark. In both studies, rigorous out-of-sample evaluation procedures are used, where possible, they use survivorship-bias-adjusted datasets, and provide several risk-adjusted performance measures, which makes cross-study comparison systematic.

The AI models under consideration belong to four different paradigms. Gradient-boosted trees (GBT) and random forests are both examples of supervised machine learning that build up ensembles of decision trees: GBT builds trees sequentially, each learning on the residuals of its predecessors (as in the XGBoost implementation), whereas random forests learn concurrently on bootstrapped subsets of the data and randomly select features to minimize variance through aggregation. Cross-sectional features (such as technical indicators (momentum at multiple horizons, RSI, Bollinger Band deviation)) and fundamental ratios (price-to-earnings, return on equity) are trained with multilayer perceptron (MLP) neural networks, using ReLU activations, batch normalization, and dropout regularization, to determine non-linear relationships between returns. In time-series deep learning, LSTM networks take input sequences of past returns and volume data; each LSTM cell can control the flow of information through input, forget and output gates, allowing the network to selectively remember or forget information over time steps, and thus learn both short-run and long-run time dependencies in financial data.

The formulation of portfolio management under reinforcement learning is a Markov Decision Process (MDP): the state contains historical prices of assets and macroeconomic factors; the action is a portfolio set of weights; and the reward is the log-portfolio net of transaction costs. Policy gradient approaches Proximal Policy Optimization (PPO) and actor-critic approaches Soft Actor-Critic (SAC) are methods to optimize allocation policies by reacting to past market environments over consecutive time steps. In evolutionary optimization, genetic algorithms encode portfolios through chromosomes of binary vectors of stock selections coupled with continuous weight vectors, and use tournament selection, uniform crossover and Gaussian mutation to solve cardinality-constrained, multi-objective optimization problems that are computationally intractable to classical quadratic programming. To standardize performance assessment across reviewed studies, annualized Sharpe ratio, Sortino ratio, maximum drawdown, and Conditional Value-at-Risk (CVaR) at the 95th percentile are used to standardize performance assessment.

Results: "The facts "

The empirical literature is consistently pointing in the same direction; AI-based portfolio strategies are better than their classical equivalents, and this performance improvement is strong across asset classes, time horizons, and methodologies. The main studies reviewed and their key performance results are summarized in table 1 with some example classical benchmark results based on DeMiguel et al. (2009) and Jorion (2007) as reference points.

Table 1: AI vs. Classical Portfolio Strategies — Empirical Evidence from Published Studies

Source / Study	AI Method	Dataset / Period	Key Metric	vs. Benchmark
Gu et al. (2020)	Neural Networks, GBT, Random Forests	~30,000 US stocks, 1957–2016	Higher out-of-sample R^2 ; NN alpha roughly doubles OLS regression	+~100% vs. OLS strategies
Fischer & Krauss (2018)	LSTM (2-layer, sequence learning)	S&P 500, 1992–2015	Sharpe ratio 5.8; daily return 0.46% (pre-cost)	Beats RF, DNN, Logistic Reg.
Jiang et al. (2017)	Deep RL — EIIE / CNN / RNN / LSTM	Cryptocurrency, 2016	Higher portfolio value growth and Sharpe ratio vs. all baselines	Beats PAMR, ONS, UBH
Chang et al. (2000)	Genetic Algorithm (cardinality-constrained)	UK, Hang Seng, DAX, FTSE equities	Near-optimal portfolios under constraints intractable for QP	Feasible vs. NP-hard QP
DeMiguel et al. (2009) [Benchmark]	Mean-Variance Optimization (14 variants)	7 empirical datasets, US equities	None of 14 MV models consistently beats 1/N in out-of-sample Sharpe ratio	Underperforms 1/N

Note: GBT = Gradient-Boosted Trees; RF = Random Forest; RL = Reinforcement Learning; EIIE = Ensemble of Identical Independent Evaluators; PAMR = Passive Aggressive Mean Reversion; ONS = Online Newton Step; UBH = Uniform Buy-and-Hold; QP = Quadratic Programming.

The most important discovery in the literature on return-prediction is the one by Gu et al. (2020). Their out-of-sample analysis reveals that neural networks can be used to approximately increase the predictive alpha of the optimal regression-based strategies by learning non-linear interactions among momentum, liquidity, and volatility signals - a phenomenon that persists after controlling for standard risk factors such as the Fama-French five-factor model (Fama and French, 2015). The fact that the dominant predictive signals are consistent across all classes of models supports the argument that AI methods are better at extracting signals, and not accidentally follow in-sample trends. Notably, even simple linear models like elastic net significantly outperform simple least squares models, showing that the advantages of machine learning are not exclusive

to complex non-linear models but also to simpler models including penalization and dimensionality reduction.

The most rigorous time-series evidence is presented by Fisher and Krauss (2018) who apply LSTM networks to the entire universe of S&P 500. The Sharpe ratio of 5.8 before costs and 0.46 percent per day are significantly higher than the random forest benchmark (Sharpe ratio is about 4.1) and the logistic regression benchmark (about 3.5). Attribution analysis shows that stocks that the LSTM network prefers are highly volatile in the near term, with reversals in the stock returns, indicating that the network is conditioned to capitalize on a well-documented but operationally difficult market anomaly. The deterioration in excess returns after 2010, frequently used as a sign that AI strategies are eventually arbitrated away, also points to a crucial strategic fact: the sustainability of AI-based alpha relies on access to proprietary information, architectural innovation and retraining-on-the-fly, not the unchanging re-instantiation of published models.

Jiang et al. (2017) RL framework can be seen as a complementary contribution. The EIIE agent can learn to maximize net-of-cost portfolio growth with a transaction-cost-penalized log-reward through framing portfolio rebalancing as an MDP, without necessarily taking an explicit return-prediction intermediate step. Backtesting findings report an improvement in performance on both Universal Portfolio and PAMR benchmarks, showing that deep RL agents are able to find dynamic allocation policies that are structurally inaccessible to a static mean-variance optimization. The genetic algorithm structure of Chang et al. (2000) further generalizes the practical application of AI-based portfolio construction to cardinality-constrained optimization problems, including restricting the portfolio to a fixed number of holdings or having a minimum position size, which makes classical quadratic programming computationally infeasible on an institutional level.

Discussion and conclusion: "The so what ? "

The combination of the empirical findings of Gu et al. (2020), Fischer and Krauss (2018), and Jiang et al. (2017) forms a logical and complementary argument in favor of the superiority of AI-based portfolio management over classical tools across three structural dimensions. First, AI methods are model-free in the sense that, instead of making parametric distributional assumptions about the returns on assets, machine learning and deep learning algorithms directly learn predictive relationships on observed data, dynamically adapting to non-linear, non-stationary dynamics that make classical mean-variance optimization unreliable out-of-sample. This adaptability is directly related to the DeMiguel et al. (2009) observation that classical optimization needs implausibly long historical samples before it can beat naive diversification, since AI methods do not need to make accurate high-dimensional covariance matrix estimates. Second, RL frameworks are a collapse of the traditional two-step quantitative pipeline of return prediction-turning into portfolio-optimization into one combined policy that directly maximizes risk-adjusted net-of-cost returns. This removes the compounding estimation errors of sequential pipeline architectures and produces agents that dynamically adjust the frequency and size of position rebates to changing market conditions. Third, genetic algorithms and the associated evolutionary strategies effectively solve the category of constrained portfolio optimization problems such as cardinality constraints, sector exposure constraints, and minimum holding constraints that is computationally infeasible to tackle using classical quadratic programming at an institutional level.

To institutional investors, these results imply the systematic use of AI-enhanced investment procedures, under significant organizational conditioning. The effective management of AI portfolios requires infrastructure to support high-quality data sourcing and feature engineering, model training, retraining, and monitoring, and a risk governance framework optimally designed to support AI systems as opposed to adapted from the practices of traditional portfolio management. The fact that LSTM excess returns have since 2010 become less pronounced, as observed by Fischer and Krauss (2018), is also a valuable reminder: AI strategies are not a fixed source of alpha: institutions should constantly innovate with proprietary data, more advanced architectures, and active retraining plans in order to ensure informational advantage. Retail investors can also more readily use AI in the form of robo-advisors and algorithmic exchange-traded funds, but the decadence in performance that Fischer and Krauss (2018) observed reminds investors against naive extrapolation of historical AI performance in future returns.

This paper has a number of limitations to its conclusions. First, the use of published research presents a selection bias since studies that report a negative result of AI are less likely to be published (Harvey, Liu, and Zhu, 2016), which may exaggerate the reported performance benefits of AI strategies. Second, the findings of RL and genetic algorithms considered are largely on cryptocurrency or simulated equity environments; their applicability to large, liquid institutional equity markets under strict regulatory supervision has not been completely determined. Third, the assumptions used in respect to transaction costs throughout the literature reviewed are simplistic and probably understate the real cost burden imposed on institutional investors, indicating that the net-of-cost performance advantages might be exaggerated. Such restrictions indicate positive research directions in the future: longitudinal trials of live-trading AI strategies, the use of multiple-testing corrections to reported AI alpha estimates, and the creation of explainability systems - including the SHAP (SHapley Additive exPlanations) analysis - that can make AI-generated portfolio choices understandable to regulators and customers. The systemic consequences of large-scale AI implementation in investment management, such as the possibility of algorithmic rebalancing to create co-ordinated patterns of trading that enhance market volatility and systemic correlation, should be of particular concern to both academic researchers and regulators.

Overall, the paper has shown that AI is a real, empirically proven paradigm shift in portfolio optimization and risk management. The overall convergence of the findings of Gu et al. (2020), Fischer and Krauss (2018), and Jiang et al. (2017) - supervised machine learning to predict returns, deep time-series forecasting, and reinforcement learning to make dynamic allocation - points to the overall conclusion that AI-based strategies can produce meaningfully better risk-adjusted returns and tail-risk management in comparison to classical. The fact that DeMiguel et al. (2009) results have shown that no classical variant of mean-variance optimization can be uniformly more successful than naive diversification highlights that the potential of quantitative investment management cannot be fully realized using only classical techniques, but requires the informational and computational power of modern AI. To achieve this potential at an institutional scale will need long-term investment in technology infrastructure, governance innovation, and regulatory engagement, which is not only technical but fundamentally organizational, and will define the competitive fabric of investment management in the coming years.

References

- Athari, S. A., Sassine, M. E., Agyemang, E. T., Kirikkaleli, D., & Saliba, C. (2026). Nonlinear Dynamics of AI Investment and Economic Growth in Germany: Evidence from a NARDL Approach. *Economies*, 14(1), 31.
- Athari, S. A., Kirikkaleli, D., Saliba, C., & Olanrewaju, V. O. (2025). Unlocking the investment nexus between artificial intelligence and Bitcoin: Modified cross-quantile regression insights. *Social Sciences & Humanities Open*, 12, 102131.
- Beainy, R. (2026). Volatility and Returns in VUCA Transatlantic Markets. *Available at SSRN 6035355*.
- Beainy, R., & Kamel, C. (2025). USA versus China From an Equity Perspective. *Available at SSRN 5742022*.
- Beainy, R. H., & Kamel, J. K. (2023). Surviving the Post Silicon Valley Bank Crisis. *International Journal of Membrane Science and Technology*, 10(3), 775-785.
- Chang, T. J., Meade, N., Beasley, J. E., & Sharaiha, Y. M. (2000). Heuristics for cardinality constrained portfolio optimization. *Computers & Operations Research*, 27(13), 1271–1302. [https://doi.org/10.1016/S0305-0548\(99\)00074-X](https://doi.org/10.1016/S0305-0548(99)00074-X)
- Kamel, C., Khalil, F., Mouawad, J., Kechli, W., & Kaspard, J. (2026). Enabling Green Innovation in the Circular Economy: A Systematic Thematic Review of Digitalization and Stakeholder Engagement. *Sustainability*, 18(3), 1360.
- Khalife, D., Yammine, J., Rahal, S., & Freiha, S. (2023). Pricing Asian and Barrier Options Using a Combined Heston Model and Monte Carlo Simulation Approach with Artificial Intelligence. *Mathematical Modelling of Engineering Problems*, 10(5).
- Khalife, D., Subrahmanyam, S., & Farah, A. (2024). A Sustainable Circular Business Model to Improve the Performance of Small and Medium-sized Enterprises Using Blockchain Technology. *J. Wirel. Mob. Networks Ubiquitous Comput. Dependable Appl.*, 15(2), 240-250.
- DeMiguel, V., Garlappi, L., & Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *Review of Financial Studies*, 22(5), 1915–1953. <https://doi.org/10.1093/rfs/hhm075>
- 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)
- Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116(1), 1–22. <https://doi.org/10.1016/j.jfineco.2014.10.010>
- Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, 270(2), 654–669. <https://doi.org/10.1016/j.ejor.2017.11.054>
- Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *Review of Financial Studies*, 33(5), 2223–2273. <https://doi.org/10.1093/rfs/hhaa009>

- Harvey, C. R., Liu, Y., & Zhu, H. (2016). ...and the cross-section of expected returns. *Review of Financial Studies*, 29(1), 5–68. <https://doi.org/10.1093/rfs/hhv059>
- Jiang, Z., Xu, D., & Liang, J. (2017). A deep reinforcement learning framework for the financial portfolio management problem. *arXiv preprint arXiv:1706.10059*.
- Jorion, P. (2007). *Value at Risk: The New Benchmark for Managing Financial Risk* (3rd ed.). McGraw-Hill.
- Ledoit, O., & Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2), 365–411. [https://doi.org/10.1016/S0047-259X\(03\)00096-4](https://doi.org/10.1016/S0047-259X(03)00096-4)
- Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, 7(1), 77–91. <https://doi.org/10.2307/2975974>
- Ross, S. A. (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, 13(3), 341–360. [https://doi.org/10.1016/0022-0531\(76\)90046-6](https://doi.org/10.1016/0022-0531(76)90046-6)
- Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19(3), 425–442. <https://doi.org/10.2307/2977928>