Optimizing Algorithmic Trading Strategies Using AI

By Dr. Vincent Wong

Associate Professor of Computer Science, Hong Kong University of Science and Technology (HKUST)

1. Introduction to Algorithmic Trading and AI

Trading involves making decisions based on predictions and achieving the highest possible profits. Algorithmic trading is a decision-making mechanism in the trading process that allows executing trades automatically based on predefined conditions. It is calculated with the help of computer algorithms. This means that computers make decisions to trade or invest. AI in trading applications is used to enhance trading strategies and replace those that require human decision-making due to their slow speed. AI can improve trading decisions when compared to those made by humans. With AI, trading decisions are less emotional, more logical, faster, fully automatic, and more optimal. It is no wonder that trading applications are developing today.

Artificial intelligence has increased market efficiency by processing larger amounts of data at higher speeds to quickly make decisions to buy or sell securities. Artificial intelligence is generally used for trading in the stock market. Algorithmic trading today is a phenomenon that is growing at all levels of financial markets. Retail and professional trading are affected by this trend and, as a result, algorithmic trading generates a large volume of trading in financial markets. Today's algorithmic trading for computer applications covers a wide range of financial markets. The application of AI technologies within the domain of algorithmic trading, including the use of machine learning algorithms, backtesting, signal generation, automated statistical arbitrage, and trade execution algorithms, is replacing human decision-making. The profound effects of AI and machine learning algorithms are very apparent in the finance industry.

1.1. Definition and Basics of Algorithmic Trading

Algorithmic trading refers to the methodology of using algorithms or sets of rules to execute trading orders in financial markets. Such orders are typically computed by computers. Strategies used in algorithmic trading are engineered in different programming languages such as R, Python, and C++. An algorithm is created by following a set of simple steps in a specific order that are meant to solve a problem and provide a solution, making the behavior of these systems easier to analyze. This increased computational power has enabled traders to use algorithms that can fully or partially automate their trading process. Automated trading strategies can offer numerous advantages compared to discretionary trading. First, they automate the execution of trades, making the process faster and more accurate. They also allow investors to formalize their investment strategy and perform backtests to verify its effectiveness. It is important to note that the popularity of algorithms and automation has transformed the markets fundamentally. A massive amount of equity trades are executed exclusively by means of sophisticated algorithms.

There is a range of strategies utilized within the algorithmic trading systems, for instance, market making, arbitrage, and execution strategies. The intention of implementing an automated trading strategy should be to trade better, in a disciplined way, and preferably to have a competitive advantage over time. In fact, only a few of the existing trading strategies are successful. Rather than automated trading being solely responsible for the ineffectiveness of trading strategies, the lack of thorough research and inefficient testing, via benchmarking and backtests, is another main cause for a strategy's subpar performance. Moreover, it is crucial for backtests and trading strategies to be gauged in financial logic where risk management effectively limits the use of leverage and investment only in liquid assets. Backtesting ensures the integrity of the trading strategy, quantifying its expected performance based on past data. Critical strategies in evaluating a backtest are risk management and their effect on perspective performance, centering on expected risk and worst-case scenario risk.

1.2. Overview of AI and Machine Learning in Finance

Artificial intelligence (AI) and machine learning are two of the key technologies transforming the financial industry. AI develops systems that can perform intelligent tasks, such as learning, reasoning, and learning from experience, while machine learning allows systems to learn patterns from available data and make decisions. These technologies help finance professionals discover profitable trading strategies, assess risks during due diligence, improve trading execution, and inform decision-making while choosing the best short-term and long-term investments. There are multiple applications in finance and investment, such as personal finance and private wealth, fraud detection, credit underwriting, algorithmic trading, robo-advisors, and customer service chatbots. While there are countless AI and machine learning models and algorithms, the most popular include artificial neural networks, which are used to model the relationships between inputs and outputs, decision trees, which make decisions based on the observed data, support vector machines, which divide values into categories based on a function, and k-nearest neighbors, which classify similar data points into clusters based on the similarity of existing data.

The key attraction to machine learning is its ability to handle large amounts of big data. For instance, in credit underwriting, machine learning algorithms can sift through millions of records and discover learned links between loan applicant attributes and historical loan default rates to help identify who is most likely to default, or in equity market prediction, machine learning algorithms can parse through vast amounts of assets and trades to unearth recurring trends in prices and seize the rare opportunities presented with positive expected return. Despite these advantages, machine learning is not without challenges. Limitations in data can lead the model to inaccurately learn relationships or cause the model to reflect biases in historical data. Furthermore, some models are beneficial in that they may offer little transparency in the reasoning and decision-making process of the model's predictions, making pump-and-dump schemes more feasible than with traditional quantitative trading strategies.

2. High-Frequency Trading Strategies

High-frequency trading (HFT) strategies have become increasingly common, accounting for 40-50% of daily stock trades in most developed stock markets. HFT strategies are characterized by several features, including a focus on very short time horizons, high order-to-trade ratios, and rapid use of powerful computers to execute numerous trades. One estimate suggested that the median latency for institutional orders from order submission to acknowledgment from the stock exchange was approximately 10 milliseconds. An HFT trader's time horizon, however, is typically on the order of 100 microseconds. The combination of these two characteristics—the high volume of transactions and substantial spending on technology—suggests that HFT traders are overreliant customers, yet a growing literature shows that HFT traders are increasingly profitable.

Three main types of HFT strategies can generally be identified: market making, statistical arbitrage, and trend following. Before exploring these in more detail, it is first necessary to distinguish between the various types of trades executed. First, a market order is one that specifies a sell or purchase of one or more given quantities of a good or instrument with the requirement that the trade be executed immediately at the present market price. Limit orders, on the other hand, are designed to trade at a given price or through a price limitation, and a stop order is a trade instruction to a security dealer designed for the broker to buy or sell, exclusively in the case that the security reaches a present or better price. HFT traders typically employ a market-making strategy that specializes as a liquidity supplier, simultaneously broadcasting quotes at which the firm is ready to sell as well as to buy stock with the intention to gain the bid/ask spread. The bid/ask spread, although on the decline today, remains an important factor allowing stock traders to recoup part of this substantial block of liquidity. High-frequency traders can quote bid-ask spreads of just one cent between $50 and $100 stocks as the observed speed of light ranges from 40 to 60 milliseconds. Despite what has been purported in the media, displaying narrow two-sided markets might enable a price impact cost reduction relative to exchanges where only the best bid and offer are shown. Finally, high-frequency traders can quote two-sided markets because of the reduction of temporary price movement risk that results from displaying post-trade information.

2.1. Definition and Characteristics of High-Frequency Trading

High-frequency trading (HFT) became ubiquitous in global financial markets over the past decade and imposed challenges in risk and security management at large. The complexity of natural phenomena potentiates algorithmic trading strategies to produce more accurate behavior at more detailed levels. Also, the smooth human interpretation is defined by three basic characteristics: (1) Rapidity, since the overall process of selecting the market and making the decision up to the actual execution of trade is done in tens of milliseconds. (2) The decision-making reliability rises from a large set of historical and real-time data and enables complex signal processing. (3) The decision-making complexity involves advanced statistical and mathematical routines fitting with data properties.

In some observations, high-frequency algorithms may use co-location facilities, specialized algorithms, and low-latency trading technology. Traders must be able to quickly locate new orders and decide what to do with them. Most of these trades are executed on dark pools, where minimal financial information is revealed to allow for such activities, primarily to screen the traders' orders from certain market participants until they are executed. Traders conduct their trades through the use of smart order routers—order routing systems that source liquidity according to the trader's specifications, sophisticated order types, and algorithms that are programmed to minimize market impact, achieve better-than-market pricing, speedy execution, and statistical arbitrage—and employ market-making and hedging algorithms to make a market in a security. Most high-frequency traders engage in statistical arbitrage as a principal strategy, that is, capturing profits from the disparities in prices of closely related products in different markets. High-frequency trading has a potential impact on market volatility as well. The traders are aggressively working to take advantage of the unusual market behavior by placing short positions because these trades require higher speeds and more complex algorithms to manage.

2.2. Advantages and Challenges of High-Frequency Trading

Many studies suggest that HFT helps to provide greater market liquidity and reduce the size of bid-ask spreads, which facilitates block trading. They help provide a fair reflection of the volume-weighted average price and use services such as continuous clearing, thereby reducing transaction costs. They are also recognized as inventory providers; in dispersing large institutional orders, they aim to reduce the level of remaining inventory and ultimately the associated price impact. It is not surprising that, without them explicitly engaging in lit trading, exchanges describe them as essential for price discovery. Moreover, we argue that algorithmic trading has indeed improved the efficiency of price discovery through the entry of new e-traders, as evidenced by both the shorter times for congestion and the increased efficacy of discovered prices more generally.

The primary potential disadvantage of HFT pertains to excessive order flow and, more generally, TSER; the former may increase "quote stuffing" or "futile trades," both of which increase volatility and cause "equity market fragility," as well as potentially reducing efficiency and price accuracy through strategies. Critics also suggest that there are potential hazards associated with the ability of HFT to generate mistaken prices and the implications of trading on those mistakes. HFT is also widely seen as a tool that enables manipulative market conduct or reduces the capacity of regulated entities to manage, monitor, or prevent market abuse. There is also potential for systemic risk; an algorithm bug led to a two-day loss of $442 million and the bankruptcy of the firm. Furthermore, the flood of orders now poses a major challenge to the monitoring departments of exchanges, which have to monitor trades in real time and retrospectively determine the underlying trading motivation, in a shift from suspicious trading to medium or high-frequency order monitoring. In this respect, it can be argued that much of the proactive responsibility has shifted from real-time surveillance and informative data sharing to policing insider and fraudulent trading in existing networks.

3. Machine Learning Techniques for Market Prediction

The previous section provided an explanation of different types of trading strategies. That means it explained the precise task to perform in order to develop an optimal trading strategy. However, other work must be done to predict market trends. This section will describe all the machine learning tools and methodologies that are most frequently used to predict market movements that other work orders to perform the optimal strategy. Among all the methodologies and techniques, the most common technique that traders and investors use to perform market analysis is called technical analysis. This technique is based on the study and interpretation of historical data. It usually studies the data extracted from the marketplace and draws future market forecasts according to past trends. Beyond that, those machine learning models, traders, or investors can also forecast the future market trend by analyzing the future using sentiment analysis or fundamental analysis.

As previously stated, there are many models that use machine learning to predict market movements, employing basically labeled data to develop a model given an output. This is commonly called supervised learning, which comprehensively accounts for strategies managers employ in building an optimal trading strategy that focuses on predictions. Besides these, there are many techniques that obtain insights and correlations without being previously labeled by a model, and all of them are known as unsupervised learning. Employing artificial intelligence in financial markets has many advantages: this management is faster, more precise, and less risky as it can recognize numerous parameters, the points where financial gains can be most quickly and efficiently made. The most important downside of using AI in the financial world is related to overfitting risk and model robustness.

3.1. Supervised Learning Models

In the vast and diverse field of machine learning, extensive research has been dedicated to implementing models for the prediction of movement in financial markets. The algorithmic trading community holds a particularly keen interest in market prediction. Supervised learning can be considered an attempt to apprehend and apply the general cause-and-effect relationships that seem to govern markets. Consequently, this technique is exploited to learn and forecast the behavioral characteristics of financial assets. As the most common type of learning, supervised learning calls for the availability of target values in the dataset. Supervised learning models are trained to anticipate an outcome by processing input variables and their corresponding values of the targeted value. General predictive models include, but are not limited to, linear regression, support vector machines, decision trees, random forests, gradient boosting, and neural networks. These methods are commonly employed for the purpose of market prediction due to their ability to learn similarities between patterns when properly fine-tuned.

The applications of supervised learning models for trading strategies are varied and far-reaching. Some strategies exploit the high interpretability of linear regression and apply it for the construction of volatility predictions in the process of portfolio rebalancing. The maximum margin separation property of support vector machines is exploited in momentum trading. One of the most exciting aspects of using machine learning techniques is their capacity to detect complex non-linear patterns in financial time series. This has been demonstrated in studies that utilize neural networks to classify fixed income products and several others that create deep learning networks for derivatives pricing. The key characteristics and desired ways in which machine learning models can be successfully implemented in a trading strategy have been identified, and each of these indicators can be addressed with a comprehensive validation and calibration process. While a multitude of various aspects may be altered and meticulously manipulated, there exist a few shared markers that may, in part, provide guidelines. These include the introduction of pairwise ratios and relevant features for volatility and momentum forecasting techniques, which may be engineered in a myriad of ways; considerations with respect to the order curve, exchange-traded fund universe, back test metrics, alpha metrics, and return holding; alterations to each predictive model's default parameters; and various pruned feature selections. Indicative information that can make the most substantial impact on trading strategy performance may be concluded from the appearance of substantial decreases or increases in trading strategy metrics.

In addition to verifying the importance through the model fit, a thorough out-of-sample validation process should be conducted. Without this step, model results ultimately become non-generalizable to new data. Due to an extensive global search space and an unstable optimal feature selection, a small in-sample feature can result in an overfit model. In this case, overfit denotes a model so tailored to the characteristics of the in-sample data that it possesses ultimate predictive capabilities in retrospect. With a model for which this is the case, its ability to correctly anticipate the characteristics of out-of-sample data becomes distorted. Thus, the development of a relevant set of features necessitates meticulous care in consideration of what could potentially be driving returns. With these features in hand, an exhaustive process of verifying model fit must be conducted such that out-of-sample performance closely reflects in-sample securities selection. In the event that Strategy A surpasses Strategy B, this will allow for a stringent assessment that the alpha improvement between combinations of features was indeed substantial.

3.2. Unsupervised Learning Models

Unsupervised learning is a type of machine learning algorithm that attempts to find structure or patterns in a set of data without predetermined outcomes. This is markedly different from supervised learning algorithms that make predictions or inferences about behavior or outcomes. The most common types of techniques used in unsupervised learning are clustering, which groups data points that share similar characteristics, and dimensionality reduction, which seeks to find a lower-dimensional representation of data while preserving as much important information as possible. Market historical data, order books, and volatility present an opportunity to find hidden relationships or patterns that occur within the market or correlate across other data series. For example, news titles or sentiment might be found to have a relationship with volatility or market price.

The power of unsupervised learning in financial data mining is similar to that offered to many benchmarks or unclassified studies. The use of unsupervised learning can allow for exploring and testing opportunities for which a label or organized classification is not previously established or known. In addition, weaknesses among existing models and their development opportunities can be revealed, including the need for new features or a new approach. An unsupervised model exposes the many relationships that are unknown to already existing supervised models, since many data scientists begin with a supervised learning approach and thus miss potential insights. The primary motivation behind a typical unsupervised model in a financial context is primarily looking for market anomalies, hidden truths, or new developments from the input features. This may not be a primary motivation running most models, yet it is an important byproduct that is beneficial for the following trading strategy development step. Models take prior statistical relationships and patterns to make conclusions and analyze the strength and appropriateness of the inferences. Additionally, one of the myriad benefits of such a model is its ability to identify possible relationships or insights that may be overlooked by constructing various supervised regression, classification, and deep learning networks. When it comes to unsupervised models, however, it should be noted that because participants are the bedrock of capital markets, variability cannot be entirely extricated. Consumers are inherently scared and risk-averse, and their actions reflect it. Governments, businesses, and institutional investors are in the same boat, with some fear and desperation bound to be in tow, resulting in uncertainty that cannot be wholly eradicated.

4. Optimization Strategies for Algorithmic Trading Systems

Optimization strategies are vital in achieving competitive edges and long-term profitable success in algorithmic trading systems. Numerous techniques can be employed to optimize algorithm parameters to maximize trading performance, including maximum profit and minimum risk. Techniques include, among others, exhaustive complete testing, genetic algorithms, particle swarm optimization, tabu search, grid search, and random search. All of these optimization mechanisms play a key role in finding the best set of parameters to optimize a trading strategy. These new parameters are then injected into the trading algorithm, which helps to keep traders ahead of markets and competitors. The overall optimization of strategies is carried out by backtesting through various market conditions, trade executions, risk management behavior, and slippage and transaction costs.

One of the most pertinent challenges in finding the best solution to a problem is the trade-off between model complexity and model interpretability. Another issue in algorithms results from the risk-return trade-offs that are directly related to every single transaction. In this context, risk management strategies are required to be integrated into the optimization process because they necessarily serve to limit risk exposure by adding more constraints in the trade selection process. Risk can be managed through increasing the stop-loss value and modifying the size of positions entered. It is also important to manage trading strategies according to a specific kind of risk and desired benefit using a combination of different strategies. Additionally, risk management techniques can be added to this kind of research to improve optimization performance. In the real world, trading strategies are dynamic multimodal constructions that should be managed and optimized continuously because trading strategies interact with market movements.

5. Case Studies and Applications of AI in High-Frequency Trading

In total, only a limited amount of real-life applications or case studies integrating AI are publicly known. Most of these examples are proprietary to their respective institutions, and details are scarce or based on media interviews. In this section, we provide a non-exhaustive list of notable case studies highlighting the use of AI in high-frequency trading strategies.

In the application, the ability of an AI system to spot structured behavior in real-time data is aimed. Implemented as a quantitative trading strategy, the system has shown to be able to respect latency and is directly flagged. Another use case is showcased by a market maker, with a level of sophistication rivaling HFT. The system developed uses an ensemble of classifiers to analyze a steady feed of real-time data, with manual judgment providing profitability for unexpected issues in the acquired data. A real-life application of AI in a fact-complex arbitrage strategy has proved its worth in replicating human judgment patterns in nearly real-time. Additionally, AI has been relied on in developing a trade execution system.

Finally, the end-of-day trading strategy designed is shown to be highly resistant to classic market behavioral biases. Used across a large portfolio, the new system needs time to gain proficiency. Both strategies presented still use AI exclusively for a partial segment of the main chain of processing. However, they underline how AI can be used, even in medium-frequency trading, in a discriminatory profiling role. Additionally, these cases also demonstrate the readiness of AI to deal with complex real-time transaction-related data for making trading decisions. Another use case is proposed, wherein AI is relied on to achieve an edge-capturing evolution on an intraday data sample. Evolution is periodically injected into the main model and traded via a downstream trading infrastructure. Finally, a multifactor long-short high-frequency trading portfolio strategy has been developed. Positions are taken based on an AI model, using supervised learning for feature extraction and learning the spread between factors. The model has an adjustable training window, and it trades a range of over 100 US equities with an average hold of a half-hour. Insight from the model has been sold as an investment product for around three years.

6. Future Direction

Markets will keep needing new strategies as they keep evolving. New investment avenues and complex financial instruments come every now and then and thus cater to the need for new strategies. The future surely holds innovative uses of technology and more advancements on the present state-of-the-art machine learning algorithms. Besides, deeper learning algorithms will also be helpful in dealing with footprints left by technical indicators and tested inefficiencies.

A faster world demands quicker changes in regulation. Regulatory authorities must act more quickly as technological advancement and its development are multiplying faster than ever. Financial authorities should keep up with technological growth and ponder the consequences of its advancement for the financial market. Finally, the strategies are not the only aspect that would leverage the usage of AI. AI will also intensify decision-making and execution processes and thoroughly improve the process of risk management.

It would be interesting to gauge the consequences of the rapid upsurge in the use of AI and how smaller players are disinclined. As AI goes mainstream, the level of competition would change. Those who cannot afford such strategies are sure to be less prosperous and would instead choose to rely on other important aspects. Hence, a different type of competition would surface. Data privacy remains a great concern. Even after all development, data could be a mode of control. This is a pressing matter that would become an issue henceforth. Lastly, any amount of mathematical prediction capabilities of AI would challenge conventional thinking as well as the acknowledgment of past paradigm shifts. It might compel us to look at the stock market in an entirely new way.

7. Conclusion

In this paper, we analyzed the dynamics of AI and algorithms and possible ways to integrate AI and machine learning to optimize algorithms used in trading financial markets. We explored the advantages of high-frequency trading as well as the challenges of it. The key objective of this work was to show that robust optimization techniques should be used to obtain trading partitions and benchmarks of algorithmic systems. The case studies presented the possibilities of integration of the most advanced and prospective techniques used for Bayesian estimation, influence diagrams, and trading systems design. Finally, we would like to present the main idea of new methodologies based on recent research. The latest trends and cutting-edge research in the field of AI can be characterized, for instance, by the trend for customization and ongoing adaptation of AI algorithms according to changes in the market. As markets change, what is profitable today may not be so tomorrow. The methodologies from these latest studies are highly adaptive and can bring benefits close to the auction dynamics of AI algorithms used for this industry.

In conclusion, this essay shows that AI is the future of the development and optimization of algorithmic trading strategies. Businesses should aim to understand key financial concepts and assess how the explosion of AI in trading techniques may affect their business practices over the coming years. Additionally, decisions related to high-frequency trading should be made using market-neutral strategies and trades using backtesting and robust optimization techniques in order to build trader splits and decision benchmarks. Until today, ethical aspects have always played an important role in trading, as they forbid using economic or market information potential. However, it is our opinion that the ethical discussion at this level of technology needs to be deepened and it needs to cover the entire trading company. It cannot be just one individual that will have a decision about whether the company should play in the future using advanced hyper-optimization and customization with AI in order to speed up its algorithms.

Reference:

1. Tamanampudi, Venkata Mohit. "Automating CI/CD Pipelines with Machine Learning Algorithms: Optimizing Build and Deployment Processes in DevOps Ecosystems." *Distributed Learning and Broad Applications in Scientific Research* 5 (2019): 810-849.
2. Pasupuleti, Vikram, et al. "Enhancing supply chain agility and sustainability through machine learning: Optimization techniques for logistics and inventory management." Logistics 8.3 (2024): 73.
3. Thota, Shashi, et al. "Federated Learning: Privacy-Preserving Collaborative Machine Learning." Distributed Learning and Broad Applications in Scientific Research 5 (2019): 168-190.
4. J. Singh, "Advancements in AI-Driven Autonomous Robotics: Leveraging Deep Learning for Real-Time Decision Making and Object Recognition", J. of Artificial Int. Research and App., vol. 3, no. 1, pp. 657–697, Apr. 2023
5. Alluri, Venkat Rama Raju, et al. "Serverless Computing for DevOps: Practical Use Cases and Performance Analysis." Distributed Learning and Broad Applications in Scientific Research 4 (2018): 158-180.
6. Machireddy, Jeshwanth Reddy. "Assessing the Impact of Medicare Broker Commissions on Enrollment Trends and Consumer Costs: A Data-Driven Analysis." *Journal of AI in Healthcare and Medicine* 2.1 (2022): 501-518.
7. S. Chitta, S. Thota, S. Manoj Yellepeddi, A. Kumar Reddy, and A. K. P. Venkata, "Multimodal Deep Learning: Integrating Vision and Language for Real-World Applications", Asian J. Multi. Res. Rev., vol. 1, no. 2, pp. 262–282, Nov. 2020
8. Ahmad, Tanzeem, et al. "Hybrid Project Management: Combining Agile and Traditional Approaches." Distributed Learning and Broad Applications in Scientific Research 4 (2018): 122-145.
9. Tamanampudi, Venkata Mohit. "CoWPE: Adaptive Context Window Adjustment in LLMs for Complex Input Queries." Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023 5.1 (2024): 438-450.

10. Thota, Shashi, et al. "Few-Shot Learning in Computer Vision: Practical Applications and Techniques." Human-Computer Interaction Perspectives 3.1 (2023): 29-59.

11. Tamanampudi, Venkata Mohit. "Leveraging Machine Learning for Dynamic Resource Allocation in DevOps: A Scalable Approach to Managing Microservices Architectures." Journal of Science & Technology 1.1 (2020): 709-748.

12. J. Singh, “Autonomous Vehicle Swarm Robotics: Real-Time Coordination Using AI for Urban Traffic and Fleet Management”, Journal of AI-Assisted Scientific Discovery, vol. 3, no. 2, pp. 1–44, Aug. 2023

13. S. Kumari, “Cloud Transformation for Mobile Products: Leveraging AI to Automate Infrastructure Management, Scalability, and Cost Efficiency”, J. Computational Intel. & Robotics, vol. 4, no. 1, pp. 130–151, Jan. 2024.