Proceedings of the 8th International Conference on Computing and Data Science
DOI: 10.54254/2755-2721/2026.36597

Artificial Intelligence in Quantitative Trading: Application Pipeline, Prospects, and Risk Governance

Zichong Long

School of Data Science and Artificial Intelligence, Guangdong University of Finance, Guangdong, China
2075623608@qq.com

Abstract. Artificial intelligence (AI) has become a significant force in the field of quantitative trading because it extends traditional rule-based systems to areas such as adaptive prediction, dynamic configuration and automated execution. At the same time, the widespread adoption of machine learning, reinforcement learning, and workflows based on large language models in financial practice has raised concerns about data overfitting, data vulnerability, herding effects, systemic risk, and governance failure. This paper provides a comprehensive review and conceptual synthesis of the application of AI in quantitative trading by integrating recent literature on machine learning, prediction and portfolio optimization in financial markets, strategy mining driven by large language models, quantitative crisis management, and AI-based risk management. This paper follows the structure of actual trading processes and explores how AI supports data collection, feature engineering, model development, portfolio construction, and execution. It further identifies four representative risk categories: model risk, data dependency and quality risk, market liquidity and volatility risk, and operational and cybersecurity risk. Based on these findings, this paper proposes a multi-layered governance framework that combines robust model engineering, investment-level controls, implementation safeguards, human oversight, and regulatory coordination. This paper argues that compared to replacing human judgment with fully autonomous algorithms, the future of quantitative trading relies more heavily on building auditable systems that combine data discipline, model validation, risk control, and human oversight.

Keywords: Artificial intelligence, Quantitative trading, Machine learning, Large language models, Risk management

1. Introduction

Quantitative trading, also known as algorithmic trading, relies on statistical models, machine learning systems, and automated execution, rather than solely on subjective judgment. In recent years, the relationship between quantitative trading and artificial intelligence (AI) has evolved from a narrow focus on technological improvement to a broader transformation that includes changes in how financial data is interpreted, how trading signals are generated, and how real-time risk management is conducted. Early quantitative strategies relied primarily on manually constructed

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/2026.36597

factors, econometric forecasts, and assumptions about the relative stability of market microstructures. Modern systems, however, increasingly rely on machine learning models capable of handling high-dimensional data, discovering nonlinear patterns, and adapting strategies to changing market conditions [1, 2].

This shift has brought both enthusiasm and concerns. On the one hand, AI has enhanced the ability of quantitative systems to analyze massive amounts of structured and unstructured information, including prices, order flows, news, market sentiment and other datasets. On the other hand, the flexibility demonstrated that AI in the financial sector has also introduced vulnerabilities. Models may overfit historical noise, inherit biases from incomplete datasets, or fail in response to changes in market conditions. Furthermore, when many institutions deploy similar models, the localized benefits of automation may accumulate into overall market instability. As a result, the question is no longer whether AI can be used for quantitative trading, but rather how it can be embedded with trading and governance frameworks to ensure that performance improvements do not come at the expense of reliability and system security [1, 3].

This paper explores three interconnected questions. First, what is the theoretical logic behind applying AI to quantitative trading? Second, how can AI be integrated into key workflows such as data collection, feature extraction, model development, and trade execution? Third, as quantitative trading systems become more automated, more data-dependent, and more adaptive, what forms of risk management are required to maintain stability and resilience? This paper aims to provide a structured perspective that links application mechanisms with risk governance, rather than treating technical performance and institutional governance as separate issues.

To answer these questions, this paper treats AI-based quantitative trading as an integrated pipeline rather than a collection of separate models. Its contribution is conceptual rather than narrowly technical. It explains how data, models, execution, and governance interact within the same trading system. By synthesizing recent research in machine learning, quantitative finance and risk governance, the paper demonstrates that the next stage of quantitative trading will be shaped by hybrid intelligence. In such a system, AI is no longer merely a predictive engine. Instead, it becomes part of a broader framework for market perception, asset allocation, execution, monitoring and compliance. This perspective is highly valuable for practitioners who not only focus on model performance, but also on deployment realism, robustness and institutional responsibility.

2. Theoretical foundations and the application pipeline of AI in quantitative trading

2.1. Theoretical rationale for AI adoption

The application of AI in quantitative trading is increasing due to the mismatch between traditional rule systems and the realities of financial markets. Because of the noise and limited observability, markets are often volatile. Moreover, participants' behavior, delays in information transmission and feedback generate a loop, which reflects on the price during signal discovery and trade execution. Given the valid model assumptions and the clear relationships between variables, classical linear models shows explanatory and predictive values. However, they still have some limitations when returns are dependent on nonlinear interactions or when the changes of inputs and outputs are observed under market mechanisms. Thereby, machine learning shows great potential in these situations. Without pre-specifying the full functional form, it can estimate complex mappings, including input variables and trading-related outcomes [1, 2].

In early quantitative trading systems, price, volume and accounting variables are the main factors. Modern AI systems can integrate these data with various information, such as corporate

DOI: 10.54254/2755-2721/2026.36597

textual disclosures, news, social media signals, macroeconomic reports, etc. Although they do not generate excess returns (alpha), they provide models to identify structures which might be overlooked. The role of large language models (LLMs) will be discussed in Section 2.4. LLMs are regarded as part of model development and strategy exploration rather than a general theoretical rationale. Consequently, Kou *et al.*'s framework is well discussed as a model-development example [4]. LLMs are used to generate factors requiring further market-aware evaluation. Furthermore, with the dynamic nature, financial markets can provide support for the AI models. Factors affecting quantitative trading include policy shocks, liquidity changes, crowding effects and technological competition. A fixed rule may be efficient for a short period of time, but it may become inefficient in long-tern run. AI models incorporating online learning offer a solution for this instability. Particularly, reinforcement learning can play an important role when the trading process involves continuous decision-making. Agents are requested to act first and gain rewards or losses based on their actions. They can adapt their behavior in response to changes in the surrounding environment accordingly.

2.2. Data collection and processing

Data collection and processing can be regarded as a risk-control stage rather than a neutral technical step. Addy et al. identify data availability and data quality as conditions for applying machine learning in financial markets as important [1]. Wu [3] links weak validation and data governance to model risk in algorithmic trading. Data governance becomes even more important in backtesting and performing evaluations. Sullivan, Timmermann and White [5] show that technical trading rules can appear profitable after repeated testing, even when the apparent performance partly reflects data-snooping rather than a stable market relation. Brown *et al.* [6] make a related point regarding survivorship bias: performance studies that retain only surviving funds or assets can overstate returns and understate risk. Consequently, an AI trading pipeline should record data lineage, align timestamps, check vendor consistency, remove look-ahead channels, and document sample construction before model training begins.

From an operational perspective, AI enters quantitative trading through a data pipeline before it appears as a model. AI systems collect and standardize market prices, fundamentals, macroeconomic variables, alternative signals, and a growing volume of unstructured text. The quality of this layer sets a hard limit on what later stages can achieve. Addy *et al.* [1] identify that biased or incomplete data can weaken both algorithmic trading performance and machine-learning-based risk management, while Wu [3] states that delayed, inconsistent, or contaminated data may lead AI systems to make faulty decisions when market conditions change. More specifically, Sullivan, Timmermann and White [5] show that data-snooping bias can make technical trading rules leading to profit when repeated testing is not under well control. Brown *et al.* [6] further demonstrate that survivorship bias can distort performance evaluation as failed or disappeared observations can be excluded from the sample. These studies suggest that data quality not only affects preprocessing issues, but also results in model risk, backtest distortion and unreliable trading decisions.

2.3. Feature extraction

Through feature extraction, raw information can be translated into representations to be adopted by models. Even at the end-to-end learning level, it remains dominant as financial signals often need economic interpretation before being used for supporting investment decisions. Although traditional

DOI: 10.54254/2755-2721/2026.36597

factors like momentum, value, volatility, and liquidity are playing an important role, AI can enrich these factors through interaction terms, latent embeddings, textual summaries and event-based representations. Ta, Liu and Addis [2] show that adding technical indicators beyond raw price and volume data can improve short-term stock prediction and support stronger portfolio outcomes. The implication is not that every available variable should be included. The more defensible lesson is that economically meaningful features, selected carefully and regularized properly, provide AI systems with a more reliable foundation.

2.4. Model development

Model development is the process of transforming processed signals into trading decisions. The choice of model depends on the specific task. Supervised learning can support profit prediction, classification methods can describe trading states, unsupervised learning can help identify trading mechanisms, reinforcement learning can model sequential decisions, and LLM-based systems can support policy exploration. In this role, LLMs can generate factor candidates and strategic hypotheses, but these outputs must still pass empirical testing, code review, and market-aware evaluation before they can influence allocation decisions [6]. No single model category is absolutely superior. When interpretability and stable out-of-sample behavior are critical, linear models may be the preferred choice. Tree-based and ensemble methods can usually strike a balance between flexibility and interpretability, while neural networks and sequence models can capture complex temporal patterns, but require more rigorous overfit control and stability testing [1, 2].

2.5. Trade execution and portfolio control

Trade execution and portfolio control determine whether predictive signals can be translated into executable strategies. Prediction alone is not sufficient. The signal must be transformed into a specific position while simultaneously accounting for transaction costs, market impact, capacity, turnover limits, liquidity constraints and risk budget. Portfolio optimization therefore remains crucial because even when a predicted signal is valuable, the performance may be poor if asset allocation is unstable, excessively concentrated, or associated with high transaction costs. Existing research from the perspective of portfolio optimization shows that the value of AI in quantitative trading depends not only on the accuracy of the predicted direction but also on how the predicted signal is transformed into an executable portfolio [2]. The practical value of AI in this stage therefore depends on whether prediction, allocation and execution are evaluated as one connected process.

3. Typical risks of AI in quantitative trading

3.1. Model risk: overfitting and algorithmic failure

Model risk is one of the core risks in AI-driven quantitative trading. It includes overfitting, model instability, and algorithm failure. AI models are powerful because they can capture complex patterns, but this flexibility also makes them prone to learning historical noise rather than persistent market patterns. Addy *et al.* [1] point out that overfitting and interpretability are key issues for machine learning-driven trading systems. Wu [3] also believes that the likelihood of algorithm failure will increase significantly when the model encounters structural fractures, policy shocks, or sudden changes in volatility and liquidity conditions . In quantitative trading, the assessment of

DOI: 10.54254/2755-2721/2026.36597

model risk depends on more than just statistical accuracy. Model evaluation also includes realized profit and loss after deducting costs, drawdown behavior, tail risk exposure, and response to stress.

Model risk can also appear before a strategy is deployed, especially when historical performance is evaluated using biased data. Sullivan, Timmermann and White [5] show that repeated testing of technical trading rules can make weak strategies look profitable when data-snooping bias is not controlled. Brown *et al.* [6] make a related point from the perspective of survivorship bias: when failed or disappeared observations are excluded, historical performance can be overstated. For AI-based quantitative trading, this issue is especially relevant. Larger feature spaces, repeated model searches, and flexible validation choices all increase the likelihood of mistaking historical coincidence for a durable trading signal.

3.2. Data dependency and quality risk

Closely connected to model risk are data dependency and quality risk. AI systems are only as reliable as the data they process. Financial datasets often include missing observations, vendor inconsistencies, revised fundamentals, survivorship bias, and timestamp mismatches. Alternative data can create additional complications because its collection processes are often noisy, legally constrained, and uneven across sectors or regions. When an AI system treats poor-quality data as a genuine signal, it may generate trades that appear internally consistent but have little economic meaning. Wu [3] emphasizes that incomplete, biased, or delayed data can mislead AI models and trigger faulty decisions. Addy *et al.* [1] also note that data bias and weak interpretability may weaken the reliability of machine-learning-based risk management itself. Brown *et al.* [6] further show that survivorship bias can distort performance evaluation when failed or disappeared observations are excluded from the sample. The rise of large language models adds another layer to this problem. LLMs can summarize disclosures, generate factors, and propose trading hypotheses, but they can also produce reasoning that appears plausible without being empirically grounded. In the framework proposed by Kou *et al.* [4], this risk is partly reduced through agent-based evaluation and category-balanced filtering. However, the need for filtering already shows that LLM outputs cannot be accepted at face value. In quantitative research, LLM-generated ideas should be treated as hypotheses rather than verified knowledge.

3.3. Market liquidity and volatility risk

Risk also comes from market liquidity and volatility risk. AI models can be optimized with smooth execution and stable market liquidity, but these conditions are not always met in the real market. When the confidence level of models reaches its highest, liquidity could become negligible and bid-ask spreads may further widen. As a result, execution costs will rise. When most institutions rely on similar optimization methods, local pursuit of excess returns will prevent the market from widening. It will lead to price volatility amplified, liquidity gaps deepened as well as self-reinforcing feedback loops formed [1, 3]. This issue is particularly severe for strategies exploiting transient microstructural effects. Actual performance can be affected by delays, queue positions and order book imbalances.

3.4. Operational and cybersecurity risk

It should draw attention to operational and cybersecurity risks. AI-based quantitative trading is highly dependent on software infrastructure, model deployment processes, data sources, cloud

DOI: 10.54254/2755-2721/2026.36597

services and execution interfaces. Various factors, such as code regression, silent data corruption, unstable dependencies, and inconsistencies between research and production environments can influence its performance. Cybersecurity will bring extra risks, including direct financial losses or serious reputational damage [3], if input data, model parameters or execution permissions were attacked. Moreover, LLM-based systems also creates other challenges, such as prompt injections, use of insecure tools, and accidental disclosure of sensitive strategy information. Therefore, AI quantitative trading platforms should be well designed to meet the requirement of software specifications, observability, access control and clear incident response processes.

4. Risk management recommendations

4.1. Multilayer defense system

Because AI-based quantitative trading creates risk at several layers, governance must be multilayered as well. The goal is defence in depth: model validation, portfolio limits, execution controls, operational safeguards, and external governance should support one another rather than function as separate checklists. This multilayer defense system mainly addresses the interaction between model risk discussed in Section 3.1 and operational risk discussed in Section 3.4. A reliable trading process should separate research validation from production deployment. Out-of-sample testing, cross-validation considering market mechanisms, evaluation incorporating transaction costs, and stress testing should all be routine practices, especially when the strategy relies on volatile signals. Backtesting is an effective screening tool, but it cannot prove future performance. Every substantive strategy should include clearly defined failure hypotheses: which market conditions might cause signal failure, which assumptions are more vulnerable, and which indicators can reveal strategy degradation early on.

Portfolio and execution safeguards should likewise be integrated within the strategy, not implemented after it. Position limits, turnover budgets, industry caps, liquidity filtering mechanisms, stop-loss rules, and circuit breakers help prevent the model's raw output from evolving into uncontrolled trading. When market conditions exceed preset thresholds, the system should scale back, tighten limits, or suspend automated trading. Manual intervention is only practically meaningful when it is technically feasible, procedurally clear, and can be initiated promptly before losses become irreversible.

4.2. Technical model robustness

Technical robustness directly responds to the overfitting and algorithmic-failure problems discussed in Section 3.1, while also reducing sensitivity to noisy or unstable data issues described in Section 3.2. Model robustness begins with restraint. If a simpler model provides comparable out-of-sample utility, it is likely preferable to a less transparent system. Feature selection, regularization, diversity in ensemble models, uncertainty estimation, and sensitivity analysis should be integrated into the model development process, rather than added at the end. Therefore, it is important for researchers and risk managers to understand the factors driving signals. Particularly the decision rules may change and economic structure or chance correlations will be reflected in the models.

Considering audit trails, LLM-assisted quantitative trading systems needs to be cautious at the same level. LLM outputs can be viewed as provisional strategies rather than approved ones. Generated factors, such as strategy rules and market narratives need to pass numerical verification, code review and performance evaluation first. Then they can be used to influence capital allocation.

DOI: 10.54254/2755-2721/2026.36597

Kou *et al.* [4] proposed a framework which adds screening and market-aware evaluation before strategy generation and after implementation. It is suggested that institutions should follow this by recording prompt templates, tool permissions, data sources and validation results.

4.3. Operational security safeguards

As discussed in Section 3.4, operational safeguards associated with the deployment and cybersecurity risks have been explained. It will be helpful to discuss further on liquidity-related failures when worse execution conditions are considered. Thereby, AI trading platforms are suggested to be managed by following production system standards. Basic requirements include version control, deployment review, access isolation, logging, rollback mechanisms and incident drills, etc. Backup plans are necessary and critical model outputs are subject to anomaly monitoring. AI systems with tool invocation capabilities should be placed in a sandbox environment. They should also be granted the necessary permissions to complete the task. From a cybersecurity perspective, real threats include data poisoning, unauthorized access and prompt injection.

4.4. Regulation and institutional development

The governance mechanism should clearly define a few factors, such as the model approver, the individual responsible for post-deployment monitoring, and the personnel authorized to report major issues or terminate transactions. Regulatory agencies can help maintain market integrity through a series of actions, including document standards, auditability of AI-assisted workflows, data source requirements, and stress testing requirements before large-scale deployment [3, 7]. This is for matching technological experiments with accountability constraints rather than preventing AI-based transactions.

Organizational learning should be incorporated into the same control system. Feedback mechanisms should be adopted by researchers, engineers, portfolio managers and risk officers. Models often fail when theoretical signals conflict with real market conditions. A healthy AI quantitative culture encourages questioning, records post-hoc analysis, and acknowledges. Strong historical evidence can also fail under new conditions. Therefore, risk management is not only for control purposes, but it also demonstrates the institution's understanding of the model.

5. Conclusion

AI is reshaping quantitative trading by expanding data layers, enriching feature construction, supporting flexible modeling, and combining prediction with portfolio control. Its contribution is reflected more in the system architecture level than in any predictive capability. However, systems that can discover complex patterns may also amplify hidden data problems and risks in models, markets, and operations, such as overfitting, crowding effects, execution costs, and operational errors. The core conclusion of the literature review is that the deployment of AI must be supported by rigorous validation and effective governance. As Quant4.0 and LLM-assisted strategy research have shown, future quantitative trading may move toward more automated, interpretable, and knowledge-driven systems. However, governance and institutional design must continue to place auditability, data sources, stress testing, and clear accountability at the core. In this sense, the long-term value of AI in quantitative trading lies not in full automation itself, but in whether human-AI trading systems can combine predictive capacity with transparency, accountability, and market resilience.

DOI: 10.54254/2755-2721/2026.36597

References

[1] Addy, W. A., Ajayi-Nifise, A. O., Bello, B. G., Tula, S. T., Odeyemi, O., & Falaiye, T. (2024). Machine learning in financial markets: A critical review of algorithmic trading and risk management. *International Journal of Science and Research Archive, 11*(1), 1853-1862.

[2] Ta, V. D., Liu, C. M., & Addis, D. (2018, December). Prediction and portfolio optimization in quantitative trading using machine learning techniques. In *Proceedings of the 9th International Symposium on Information and Communication Technology* (pp. 98-105).

[3] Wu, X. (2026). Risk management in AI-based quantitative trading: A comprehensive review and perspective. *Frontiers in Economics and Management, 7*(1), 198-201.

[4] Kou, Z., Yu, H., Luo, J., Peng, J., Li, X., Liu, C., Dai, J., Chen, L., Han, S., & Guo, Y. (2024). Automate strategy finding with LLM in quant investment. *Seed, 1*, 3.

[5] Sullivan, R., Timmermann, A., & White, H. (1999). Data-snooping, technical trading rule performance, and the bootstrap. *The Journal of Finance, 54*(5), 1647-1691.

[6] Brown, S. J., Goetzmann, W., Ibbotson, R. G., & Ross, S. A. (1992). Survivorship bias in performance studies. *The Review of Financial Studies, 5*(4), 553-580.

[7] Guo, J., Wang, S., Ni, L. M., & Shum, H. Y. (2024). Quant 4.0: Engineering quantitative investment with automated, explainable, and knowledge-driven artificial intelligence. *Frontiers of Information Technology & Electronic Engineering, 25*(11), 1421-1445.