

Intelligence-driven Quantitative Trading Strategies and Practices

Yujun Lu^{1,*}

¹School of Software, Nanjing University of Information Science and Technology, Jiangsu, China

*Corresponding author: YujunLu_Rise@outlook.com

Abstract:

Modern financial markets generate high-volume and high-speed data. Although machine learning has improved prediction and decision-making, guidelines for implementation considering costs, governance, and robustness against regime changes remain limited. This study comprehensively summarizes how supervised learning, unsupervised learning, deep learning, ensemble learning, and reinforcement learning align with the pipeline of quantitative trading in signal generation, portfolio construction, execution, and market making. Representative models, such as support vector machines, gradient boosting trees, and long short-term memory (LSTM) networks, are associated with traditional frameworks for allocation, execution, and inventory management, forming a decision-making centered approach. A structured review and a comparative mapping under realistic constraints are used. Evaluation methods to mitigate leakage and overfitting are summarized, and backtesting with pre-screening that explicitly considers trading costs and market impact is also included. The results reveal complementary strengths: ensemble trees and support vector machines provide a robust baseline, sequential models with attention improve temporal representation, reinforcement learning aligns with the goals of sequential trading but requires careful reward design and reliable offline training, and unsupervised tools reveal patterns that support diversification.

Keywords: machine learning; quantitative trading; reinforcement learning; portfolio optimization; ensemble learning

1. Introduction

Modern markets are data-rich and evolving rapidly. Machine learning provides tools for learning non-linear state-dependent models from large, noisy datasets

and for linking predictions to decision-making. Recent research has reported significant improvements in asset price setting and prediction, and reinforcement learning is rapidly gaining popularity in trading research [1,2]. These advancements are prompting a

comprehensive view of how major paradigms are applied to the quantitative trading pipeline and where the areas of sustainable value creation lie.

This study aims to address this need. In this study, it reviews supervised learning, unsupervised learning, deep learning, ensemble learning, and reinforcement learning in signal generation, portfolio construction, execution, and market trading. This study places emphasis on a decision-making centered approach. Signals need to be transformed into portfolios and orders, explicitly considering costs, risks, and constraints. This end-to-end approach is based particularly on classical frameworks such as asset allocation by the mean-variance model and its extensions, execution by Almgren-Chriss, and inventory price setting by Avellaneda-Stoikov [3-5].

There are three contributions. First, representative algorithms are summarized, including the algorithm's expansion tree, sequential models, and deep reinforcement learning, etc. Emphasis is placed on the evidenced conditions under which they are effective. Second, the discussion extracts implementation lessons often omitted in model-centered research, such as realistic costs, market impacts, capacity limitations, and governance of data lineage and reproducibility. Third, these insights are associated with evaluation practices, and emphasis is placed on the pitfalls of multiple tests and overfitting in backtesting, which may overestimate the reported performance.

Methodologically, the study unfolds in four stages. Section 2 sets up a brief technical context on machine learning and the trading stack. Section 3 examines the applications by paradigm with concise case proofs and model schemas. Section 4 discusses the strengths and limitations, organized around the robustness against non-stationarity, the alignment between learning objectives and trading objectives, and operational constraints. Section 5 concludes with guidelines for research and deployment. The selection gives preference to peer-reviewed sources or those archived in the last ten years, and, if necessary, to fundamental works.

The aim is practical: to provide clear guidelines focused on machine learning decision-making in quantitative trading. The specific objectives are to identify the types of models suitable for different stages of the pipeline, formulate learning goals taking into account costs and risks, and evaluate strategies using robust procedures for data exploration. The desired outcome is not only predictive accuracy but also a performance that is applicable and robust even out-of-sample.

2. Overview of Machine Learning and Quantitative Trading

2.1 Machine Learning

Machine learning covers supervised, unsupervised, rein-

forcement, deep, and ensemble methods, each learning from data under different information regimes [6]. To move beyond taxonomy, this section focuses on three representative models that anchor much of modern practice: support vector machines, gradient-boosted decision trees, and long short-term memory networks.

Support Vector Machine constructs a maximum-margin separator in a possibly high-dimensional feature space via kernel mappings. The primal objective balances margin width and hinge-loss penalties; the dual enables efficient use of kernel functions for non-linear decision boundaries. Soft-margin variants handle noise through a regularization constant, and multi-class extensions are built by one-vs-rest or one-vs-one decompositions [7].

Gradient boosting builds an additive model by stage-wise functional gradient descent. Scalable implementations such as XGBoost add sparsity-aware split finding and second-order optimization to improve speed and accuracy across regression, classification, and ranking tasks [8].

Long Short-Term Memory is a recurrent network with gating mechanisms. It is engineered to capture long-range dependencies and simultaneously prevent the problem of vanishing gradients. Input, output, and forget gates regulate information flow through a cell state, enabling stable credit assignment across long sequences; bidirectional stacks and backpropagation-through-time make it a standard baseline for sequence modeling [9].

2.2 Quantitative Trading

Quantitative trading automates signal generation, portfolio construction, execution, and market-making under explicit models and constraints [10]. To keep the discussion concrete, this section highlights three foundational models: mean-variance and Black-Litterman for portfolio selection, Almgren-Chriss for optimal execution, and Avellaneda-Stoikov for limit-order market making.

Mean-variance portfolio selection solves a quadratic program that trades expected return against variance under budget and, when needed, additional constraints. The Black-Litterman formulation blends market-implied equilibrium returns with investor views through a Bayesian update, stabilizing weights and improving out-of-sample behavior relative to naïve mean-variance inputs [3,11].

Optimal execution in the Almgren-Chriss framework chooses a trade schedule that minimizes a mean-variance objective given temporary and permanent price impact and exogenous volatility; closed-form trajectories arise under quadratic impact [4].

Limit-order market making under Avellaneda-Stoikov models inventory-aware quotes with Poisson order arrivals. Optimal bid and ask offsets balance spread capture against inventory and adverse-selection risk via a Hamilton-Jacobi-Bellman equation, producing dynamic quotes that adapt to inventory and market conditions [5].

Risk management overlays these modules with exposure,

leverage, and drawdown controls; Expected Shortfall is often preferred to Value-at-Risk for tail-risk sensitivity [12].

3. The Application of Machine Learning in Quantitative Trading

With the growth of data and the enhancement of computing capabilities, machine learning has found extensive applications in quantitative trading. The latest progress in machine learning and deep learning has transformed the methods of stock prediction. Unlike traditional statistics, these models learn latent patterns from large historical datasets and respond in near real time [1]. Accurate trend prediction can improve portfolio decisions and risk control. This chapter focuses on five method families and highlights representative applications and models.

3.1 Supervised Learning in Quantitative Trading

Supervised models learn from labeled market data to predict trends and generate signals. They are commonly used for direction classification and return or price regression [6]. In financial prediction tasks, stacked and hybrid classifiers, which integrate logistic regression, decision trees, support vector machines, and recurrent networks, have shown accuracy improvements compared to single models [13].

Feature engineering remains crucial. A three-stage feature pipeline-based hybrid GA-XGBoost method enhanced the prediction of the next-day direction of the KOSPI index [14]. Dimensionality reduction combined with news and macro information has also been explored to enhance stability and regime discrimination [15].

3.2 Unsupervised Learning in Quantitative Trading

Clustering groups stocks or market regimes by similarity and supports portfolio construction and risk monitoring. Recent studies use unsupervised algorithms to identify regime structure and latent market factors from returns, news, and macro indicators [15]. Clustering also assists pairs trading; unsupervised selection yields more stable pairs and trading returns than traditional methods [16].

3.3 Reinforcement Learning in Quantitative Trading

Reinforcement learning learns trading policies through interaction with the market environment and has been applied to single-asset trading, portfolio allocation, execution, and market making. Ensembles of deep Q-learning agents have shown positive cumulative returns and better stability in high-frequency settings [2]. Sentiment-aware deep reinforcement learning enhances the robustness of

portfolio allocation in large equity universes [17]. Surveys also report progress in order execution and inventory-aware market making using reinforcement learning, while noting challenges in reward design, safe exploration, and sim-to-real transfer [18].

3.4 Deep Learning in Quantitative Trading

Deep learning captures nonlinear temporal structure and heterogeneous features. For time-series prediction, LSTM networks often outperform shallow models; for low-volatility stocks, a two-layer LSTM achieved strong accuracy relative to traditional baselines [19]. Bi-directional LSTM and related architectures have been applied to short-horizon prediction with favorable results [20].

Convolutional networks operate on image-like representations such as K-line mosaics, and combinations with deep reinforcement learning have produced profitable signals in several settings [2,13]. Multimodal pipelines embed news into vectors and join them with price sequences in attention-based recurrent or transformer backbones, improving performance during volatile periods [21,22]. Recent work further explores large language models to structure financial text and support downstream predictors [23].

Generative approaches are emerging. Hybrid models that combine sentiment analysis, denoising autoencoders, and LSTM improved both predictive accuracy and risk-return metrics [21]. GAN-based hybrids with evolutionary optimization have outperformed pure GAN and traditional statistical models in controlled studies [24].

3.5 Ensemble Learning in Quantitative Trading

Ensembles improve accuracy and robustness by aggregating diverse models. Graph-temporal ensembles first extract cross-stock relations with graph convolutions and then model temporal dynamics with recurrent networks; this approach has shown stability in bear markets and superior performance to earlier baselines [25]. For multimodal fusion, separate predictors for technical, fundamental, and news signals can be combined by weighted averaging to yield higher accuracy than any constituent model [22].

4. Discussion

4.1 Strengths of Machine-Learning-based Quantitative Trading

Machine learning captures nonlinear relations, feature interactions, and regime-dependent patterns that linear factor models miss. Deep networks learn hierarchical representations directly from raw signals, and ensembles reduce variance and stabilize performance.

Modern pipelines integrate heterogeneous sources by combining prices and volumes, fundamental indicators,

alternative data, and microstructure features within a single learning system. This integration broadens usable signals and improves robustness.

Loss and reward functions can be aligned with trading goals such as risk-adjusted return, turnover control, and drawdown aversion. Reinforcement learning optimizes sequential decisions including timing, position sizing, and inventory management, rather than separating prediction from execution.

Adaptive procedures support changing markets. Online or rolling retraining, transfer or meta-learning, and principled model selection or ensembling enable continuous adjustment to shifting regimes. Automated feature generation and hyperparameter optimization accelerate research iteration and deployment.

Execution quality also improves. Supervised models and reinforcement learning learn cost-aware slicing, scheduling, and quoting policies that reduce slippage and market impact, which lowers implementation shortfall compared with heuristic execution.

4.2 Limitations of Machine-Learning-based Quantitative Trading

Financial data are non-stationary and subject to concept drift. Policy changes, structural shifts, and crowding alter the data-generating process. Signals can decay quickly, and models trained on past regimes often fail near regime breaks.

Overfitting is a persistent risk. High-dimensional features, limited effective sample sizes, and extensive hyperparameter search increase false discoveries. Typical pitfalls include look-ahead bias, leakage across time, and backtest overfitting. Without purged or embargoed walk-forward evaluation and strict research hygiene, many reported alphas are artifacts.

Predictive accuracy does not guarantee profitability. Transaction costs, financing, borrow fees, and market impact can erase paper gains. Unconstrained learners may generate excessive turnover and fragile profit and loss.

Reinforcement learning introduces additional challenges. Sample inefficiency, stochastic and delayed rewards, and unsafe exploration hinder stable learning. Gaps between simulators and live markets degrade policies, and reward design that reflects risk, capacity, and costs remains difficult.

Limited interpretability complicates governance and audit. Black-box models impede risk oversight and regulatory review. Results can be sensitive to data vintages, feature choices, and random seeds, which harms reproducibility and model risk management.

Operational constraints restrict feasible architectures. Latency budgets, compute and memory limits, and data quality issues such as survivorship bias, stale quotes, corporate actions, and microstructure noise are common. Monitoring, drift detection, and safe rollback are essential

but often under-engineered.

Markets adapt to profitable strategies. Crowding and adversarial responses compress edges, raise impact, and impose capacity limits. Successful policies can alter microstructure and invalidate their own assumptions.

5. Conclusion

This study reviewed how core machine-learning paradigms, including supervised, unsupervised, reinforcement, deep, and ensemble methods, interact with the quantitative trading pipeline across signal generation, portfolio construction, execution, and market making. The topic is important because modern markets are data-rich and fast-moving. Methods that learn nonlinear, regime-dependent structure and that connect prediction to decision can improve risk-adjusted performance. These methods also help scale research and operations.

The synthesis in Sections 2–4 clarifies where machine learning adds durable value and where it fails. Tree ensembles and support vector machines provide strong baselines with interpretable structures. LSTM and CNN hybrids with attention improve sequence modeling. Reinforcement learning aligns learning with sequential trading objectives. Unsupervised tools reveal regimes and latent factors that aid diversification. Credible deployment requires design choices that were often implicit in prior work. Purged, walk-forward evaluation with realistic costs and market impact is necessary. Objectives should account for risk, turnover, and drawdown. Capabilities and personnel allocation need to be evaluated. Governance must ensure the traceability, reproducibility, and explainability of data. Associating signals with implementation is just as important as mere predictability. Portfolio optimization, execution plans, and inventory management need to be designed together with the learning module and should not be added after the model has been formed.

Several methods can narrow the gap between promising past tests and robust online performance. Robustness to non-stationarity can be improved through online learning, meta-learning, change point detection, and domain adaptation between assets and regimes. Decision-centric learning benefits from end-to-end pipelines that associate predictions with the portfolio and execution layers, as well as learning objectives that incorporate differentiable risk and cost terms. Reliable reinforcement learning requires safer exploration, offline training on high-fidelity simulators, and real-world validation of simulators under cost and capacity constraints. Multimodal and base models require reasonable integration of prices, fundamental data, and text, strict control of information leakage, and estimation of uncertainty for position sizing. Evaluation and governance need to be based on open benchmarks that are sensitive to trading costs, stress tests for extreme events, and fully documented data vectors with reproducible artifacts.

Machine learning expands the quantitative toolbox from prediction to optimization of decision-making. The sustainable impact depends on approaches that are accurate, cost-conscious, risk-controlled, interpretable, and operationally robust.

References

- [1] Gu S., Kelly B., Xiu D. Empirical asset pricing via machine learning. *The Review of Financial Studies*, 2020, 33(5): 2223-2273.
- [2] Carta S., Ferreira A., Podda A. S., et al. Multi-DQN: An ensemble of Deep Q-learning agents for stock market forecasting. *Expert Systems with Applications*, 2021, 164: 113820.
- [3] Markowitz H. Portfolio Selection. *Journal of Finance*, 1952, 7(1): 77-91.
- [4] Almgren R., Chriss N. Optimal execution of portfolio transactions. *Journal of Risk*, 2000, 3(2): 5-39.
- [4] Almgren R., Chriss N. Optimal execution of portfolio transactions. *Journal of Risk*, 2000, 3(2): 5-39.
- [5] Avellaneda M., Stoikov S. High-frequency trading in a limit order book. *Quantitative Finance*, 2008, 8(3): 217-224.
- [6] James G., Witten D., Hastie T., Tibshirani R., et al. *An Introduction to Statistical Learning: with Applications in Python*. Cham: Springer, 2023.
- [7] Cortes C., Vapnik V. Support-vector networks. *Machine Learning*, 1995, 20(3): 273-297.
- [8] Chen T., Guestrin C. XGBoost: A Scalable Tree Boosting System. In: *KDD '16*. 2016: 785-794.
- [9] Hochreiter S., Schmidhuber J. Long Short-Term Memory. *Neural Computation*, 1997, 9(8): 1735-1780.
- [10] Chan E. P. *Quantitative Trading: How to Build Your Own Algorithmic Trading Business*. 2nd ed. Wiley, 2021.
- [11] Black F., Litterman R. Global Portfolio Optimization. *Financial Analysts Journal*, 1992, 48(5): 28-43.
- [12] Acerbi C., Tasche D. Expected Shortfall: A natural coherent alternative to Value at Risk. *Economic Notes*, 2002, 31(2): 379-388.
- [13] Kamara A. F., Chen E., Pan Z. An ensemble of a boosted hybrid of deep learning models and technical analysis for forecasting stock prices. *Information Sciences*, 2022, 594: 1-19.
- [14] Yun K. K., Yoon S. W., Won D. Prediction of stock price direction using a hybrid GA-XGBoost algorithm with a three-stage feature engineering process. *Expert Systems with Applications*, 2021, 186: 115716.
- [15] Barradas A., Canton-Croda R. M., Gibaja-Romero D. E. Identification of patterns in the stock market through unsupervised algorithms. *Analytics*, 2023, 2(3): 592-603.
- [16] Han C., He Z., Toh A. J. W. Pairs trading via unsupervised learning. *European Journal of Operational Research*, 2023, 307(2): 929-947.
- [17] Koratamaddi P., Wadhwani K., Gupta M., et al. Market sentiment-aware deep reinforcement learning approach for stock portfolio allocation. *Engineering Science and Technology, an International Journal*, 2021, 24(4): 848-859.
- [18] Alameer A., Saleh H., Alshehri K. Reinforcement learning in quantitative trading: A survey. *Authorea Preprints*, 2022.
- [19] Deepika N., Nirupama Bhat M. An efficient stock market prediction method based on Kalman filter. *Journal of The Institution of Engineers (India): Series B*, 2021, 102(4): 629-644.
- [20] Sunny M. A. I., Maswood M. M. S., Alharbi A. G. Deep learning-based stock price prediction using LSTM and bi-directional LSTM model. In: *2020 2nd Novel Intelligent and Leading Emerging Sciences Conference (NILES)*. IEEE, 2020: 87-92.
- [21] Zhao Y., Yang G. Deep learning-based integrated framework for stock price movement prediction. *Applied Soft Computing*, 2023, 133: 109921.
- [22] Jing N., Wu Z., Wang H. A hybrid model integrating deep learning with investor sentiment analysis for stock price prediction. *Expert Systems with Applications*, 2021, 178: 115019.
- [23] Cao B., Wang S., Lin X., et al. From deep learning to LLMs: a survey of AI in quantitative investment. *arXiv preprint arXiv:2503.21422*, 2025.
- [24] Polamuri S. R., Srinivas K., Mohan A. K. Multi-model generative adversarial network hybrid prediction algorithm for stock market prices prediction. *Journal of King Saud University - Computer and Information Sciences*, 2022, 34(9): 7433-7444.
- [25] Sun Z., Harit A., Cristea A. I., et al. Money: Ensemble learning for stock price movement prediction via a convolutional network with adversarial hypergraph model. *AI Open*, 2023, 4: 165-174.