IJARSCT
ISSN: 2581-9429
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
Volume 6, Issue 7, March 2026
Impact Factor: 8.2

Machine Learning for Stock Valuation and Portfolio Optimization

Shreyas S, Gagandeep C, Dr Mahesh Kumar N
Department of Information Science
Global Academy of Technology, Bangalore, Karnataka, India
shreyas1ga22is148@gmail.com, gagandeep1ga22is055@gmail.com, maheshkumar.n@gat.ac.in

Abstract: *This survey provides a systematic review of machine learning techniques applied to stock valuation and portfolio optimization, covering the period from foundational approaches through contemporary deep learning methods. We examine four principal research streams: (i) traditional valuation frameworks and their limitations, (ii) recurrent and attention-based architectures for price forecasting, (iii) reinforcement and imitation learning for automated fundamental analysis, and (iv) high-performance computing platforms for real-time portfolio construction. Across these domains, we identify prevailing methodologies, benchmark datasets, evaluation protocols, and performance trends. The survey further delineates open research challenges including cross-market generalization, data quality constraints, model interpretability, and the computational demands of tick-level processing. By synthesizing findings from over ten primary studies, this work offers researchers and practitioners a structured understanding of the current state of the field and a roadmap for future investigation.*

Keywords*:* Stock Valuation, Portfolio Optimization, Deep Learning, Transformer Models, Discounted Cash Flow, Imitation Learning, Time Series Forecasting, High-Performance Computing

I. INTRODUCTION

The intersection of machine learning and quantitative finance has produced a rapidly expanding body of literature over the past decade. Accurate stock price forecasting and efficient portfolio construction are two of the most consequential challenges in financial engineering, with direct implications for investment strategy, risk management, and systemic market stability. Whereas conventional approaches rely on analyst-driven heuristics and closed-form optimization, modern machine learning offers data-driven alternatives capable of capturing nonlinear dependencies, high-dimensional feature interactions, and temporal dynamics that elude classical models.

Despite this promise, the literature remains fragmented. Studies on price prediction often proceed independently of portfolio construction research, and fundamental valuation techniques such as Discounted Cash Flow (DCF) analysis have only recently been subjected to systematic automation via learning-based methods. Furthermore, the computational infrastructure required to support real-time decision-making at scale has received comparatively less attention in the academic literature, even as practitioners increasingly demand low-latency systems.

This survey aims to consolidate and critically assess the state of the art across these interconnected domains. We organize existing work into four thematic areas and evaluate each according to methodological rigor, empirical performance, and practical applicability. The remainder of this paper is structured as follows: Section II discusses the scope and methodology of the review. Section III surveys traditional valuation approaches. Section IV reviews machine learning methods for price prediction. Section V covers reinforcement and imitation learning for fundamental valuation. Section VI examines high-performance portfolio optimization platforms. Section VII synthesizes cross-cutting findings and identifies research gaps. Section VIII concludes.

DOI: 10.48175/IJARSCT-32220

II. SURVEY SCOPE AND METHODOLOGY

The primary literature surveyed in this work was drawn from IEEE Xplore, Procedia Computer Science, and related peer-reviewed venues, covering publications between 2020 and 2025. Search terms included combinations of 'stock price prediction', 'portfolio optimization', 'deep learning finance', 'Transformer time series', 'imitation learning valuation', and 'high-performance computing trading'. Initial screening yielded over forty candidate papers; after applying inclusion criteria requiring empirical evaluation on financial datasets, ten core studies were selected for detailed analysis.

Inclusion criteria required that each study: (a) propose or evaluate a machine learning method directly applied to stock-related tasks, (b) report quantitative performance metrics on at least one standard benchmark or real-world financial dataset, and (c) be published in a peer-reviewed conference proceeding or journal. Studies focusing exclusively on cryptocurrency, forex, or commodity markets without comparison to equity markets were excluded.

The selected studies are analyzed along the following dimensions: model architecture, input feature representation, training and evaluation protocol, primary performance metric, and key limitations acknowledged by the authors.

III. TRADITIONAL VALUATION FRAMEWORKS

3.1 Discounted Cash Flow Analysis

The Discounted Cash Flow method remains the most widely cited approach to intrinsic equity valuation and serves as the conceptual foundation for several machine learning extensions reviewed in this survey. The DCF framework computes intrinsic value by discounting projected future free cash flows (FCF) at the weighted average cost of capital (WACC), summing over a finite projection horizon and adding a terminal value component. Its appeal lies in theoretical grounding within corporate finance; its limitation lies in the sensitivity of estimates to appraiser assumptions regarding FCF growth trajectories and discount rate selection.

Malafeyev et al. [8] applied DCF methodology to the specific domain of social network valuation, surfacing challenges that arise when conventional frameworks are transposed to technology firms with atypical cash flow structures. Their work demonstrates that even the mechanics of applying DCF require domain-specific adaptation, a finding that motivates the learning-based approaches discussed in Section V. The study highlights that subjective parameter estimation introduces systematic errors that undermine prediction consistency across practitioners.

The broader implication for machine learning researchers is that traditional valuation cannot simply be bypassed; rather, its domain knowledge provides a structured inductive bias that, when properly encoded, can improve the generalization of learned models. This observation underpins the hybrid methodological stance adopted by Peng and Lee [3], discussed in Section V.

3.2 Comparative Ratios and Fundamental Indicators

Beyond DCF, ratio-based methods including price-to-earnings, price-to-book, and return on equity are widely used as screening heuristics. While individually weak predictors, these fundamentals have been incorporated as auxiliary input features in several classification and regression frameworks reviewed in subsequent sections. Their primary value in the machine learning context lies not in isolation but as conditioning variables that provide economic context alongside time-series price signals.

IV. MACHINE LEARNING METHODS FOR STOCK PRICE PREDICTION

4.1 Recurrent Architectures: RNN and LSTM

The earliest wave of deep learning applied to financial time series employed Recurrent Neural Networks (RNNs) and their gated variants. The Long Short-Term Memory (LSTM) network, introduced to address vanishing gradient problems in vanilla RNNs, became the de facto baseline for sequential financial data. Goyal and Raj [5] conducted a systematic comparative study of LSTM, Deep Neural Networks (DNN), and the classical ARIMA statistical model, evaluating each on closing price data for four major indices: NASDAQ (IXIC), Dow Jones (DJI), S&P 500 (GSPC),

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 7, March 2026
Impact Factor: 8.2

and Nikkei 225 (N225). The study reported RMSE comparisons across all model configurations, establishing that LSTM consistently achieves lower prediction error than both ARIMA and vanilla DNN architectures on these benchmarks.

Bansal et al. [9] extended this line of inquiry by evaluating five algorithms on a twelve-company stock database: K-Nearest Neighbors (KNN), linear regression, support vector regression, decision tree regression, and LSTM. Their Procedia Computer Science study confirmed LSTM's superiority in capturing temporal autocorrelation in equity returns, positioning it as a strong baseline against which subsequent Transformer-based models are appropriately measured.

4.2 Transformer-Based Architectures

The introduction of the Transformer architecture to natural language processing prompted rapid adoption in financial sequence modeling. Two representative applications reviewed here are the Transformer-Time2Vec framework and the Informer model.

4.2.1 Transformer with Time2Vec Encoding

Lee and Yoo [1] proposed a hybrid architecture combining multi-head self-attention with Time2Vec temporal encoding. Time2Vec transforms scalar time inputs into vector representations that decompose temporal variation into a linear trend component and a set of learnable sinusoidal components, enabling the model to capture both long-run progressions and periodic intraday or seasonal patterns simultaneously. The Transformer component then operates on these enriched representations, computing pairwise relevance scores across all positions in the input window via scaled dot-product attention.

Empirical evaluation on historical equity data demonstrated an R^2 of 0.92 and directional accuracy of 68.5%, outperforming LSTM ($R^2 = 0.91$) and Kalman filter baselines ($R^2 = 0.88$). The model required six encoder layers and eight attention heads, with a hidden dimension of 512, trained over 100 epochs on a single GPU in approximately 4.2 hours. The authors attribute the performance gain primarily to the Time2Vec encoding's ability to disentangle periodic market seasonality from secular price trends.

4.2.2 Informer for Long-Sequence Forecasting

Yulistiani and Kurniadi [2] evaluated the Informer model, which replaces standard full self-attention ($O(L^2)$ complexity) with a ProbSparse mechanism ($O(L \log L)$) that retains only the top-scoring query-key interactions. This architectural choice enables efficient processing of substantially longer input sequences than standard Transformers, a relevant property when modeling multi-month historical contexts.

The study observed that Informer performance is markedly dataset-dependent, with strongest results on banking sector equities and higher prediction errors on technology stocks attributed to elevated volatility. The authors suggest that the model's approximation of attention may sacrifice precision on high-noise series, and that careful per-sector hyperparameter tuning is required. This finding highlights a practical limitation for real-world deployment across heterogeneous stock universes.

4.3 Non-Parametric Approaches: Enhanced KNN

In contrast to deep learning approaches, Huang [6] proposed an improved K-Nearest Neighbors model that operates on N-day price trend curves rather than single time-point features. The method applies trend centralization to normalize series to a common reference frame and employs offset correction during inference to translate relative predictions back to absolute price levels. On the reported benchmarks, this enhanced KNN achieved an RMSE of 3.660 compared to 3.981 for traditional KNN and 4.221 for logistic regression.

DOI: 10.48175/IJARSCT-32220

While this approach offers no guarantee of scalability to high-dimensional feature spaces, it provides interpretability advantages and competitive accuracy for stocks exhibiting strong momentum properties. Its computational simplicity also makes it attractive as a real-time inference component where latency budgets are tight.

Table 1: Comparative Summary of Price Prediction Methods

Method	RMSE	R^2
Linear Regression	0.071	0.85
Kalman Filter	0.063	0.88
Vanilla RNN	0.052	0.89
LSTM	0.048	0.91
Transformer-Time2Vec [1]	0.050	0.92
Informer [2]	0.048 (banking)	—

V. REINFORCEMENT AND IMITATION LEARNING FOR FUNDAMENTAL VALUATION

5.1 Automating DCF via Imitation Learning

Peng and Lee [3] presented arguably the most technically novel contribution reviewed in this survey: a framework that automates the DCF valuation process through a combination of imitation learning and guided policy search. The central challenge they address is the subjectivity of expert FCF projection, which they reformulate as a learning problem. A teacher agent demonstrates state-action pairs derived from financial domain knowledge, and a student policy is trained to replicate these demonstrations while generalizing to unseen companies.

The FCF trajectory over a company's business lifecycle is modeled using the Lévy distribution, which captures the characteristic asymmetry of growth, maturity, and long-tail decline stages. Guided policy search then optimizes the trajectory parameters (pivot point, scale, growth rate) by minimizing the squared deviation between the modeled and observed historical FCF series. This parameterization provides a principled inductive bias that constrains the search space while remaining flexible enough to accommodate diverse business models.

WACC estimation is handled via the Capital Asset Pricing Model, with beta computed through regression of historical stock returns against the market index. The integrated system produces intrinsic value estimates without requiring analyst input, addressing the reproducibility and scalability concerns identified in Section III.

5.2 Performance and Implications

In empirical evaluation, the DCF-Imitation Learning approach achieved a Sharpe ratio of 1.34 for portfolios constructed from its top-ranked stocks, compared to 1.12 for Random Forest classification, 0.89 for P/E-based screening, and 0.68 for random selection. Precision in identifying genuinely undervalued stocks reached 71.3%, with recall of 64.7%. These results suggest that machine learning can operationalize the intuitions underlying expert fundamental analysis without sacrificing rigor, while substantially improving reproducibility.

The authors note that the Lévy distribution model is well-suited to stable or mature firms but may require extension for early-stage growth companies whose FCF trajectories deviate sharply from historical norms. This is a meaningful limitation for technology-sector applications.

Table 2: Fundamental Valuation Method Comparison

Method	Sharpe Ratio	Annual Return
Random Selection	0.68	9.2%
P/E-Based Screening	0.89	12.5%
Random Forest Classification	1.12	15.8%
DCF-Imitation Learning [3]	1.34	18.6%

VI. HIGH-PERFORMANCE PORTFOLIO OPTIMIZATION PLATFORMS

6.1 Distributed Computing Architectures

Chen et al. [7] tackled the computational bottleneck in large-scale portfolio optimization through the design of a High-Performance Portfolio Optimization (HPPO) platform. The platform is organized as a three-layer architecture: a data layer responsible for market data ingestion and cleaning, a model layer handling parameterized portfolio strategy construction and backtesting, and an execution layer managing order generation and real-time rebalancing. The system is built on a master-slave parallel topology, with the master node broadcasting periodic market snapshots to slave nodes that conduct independent parameter search in parallel.

To accommodate different data frequencies, the broadcast period is calibrated to data volume: annual periods for daily data, monthly periods for minute-level data, and weekly periods for tick-level data (approximately 860 MB per round). This design ensures that each slave node receives a complete and consistent view of the market without requiring synchronous communication at every tick.

6.2 Benchmarking Against Existing Frameworks

The HPPO platform was benchmarked against two widely used open-source backtesting frameworks, Zipline and Rqalpha, under serial and parallel computing conditions. In serial evaluation, HPPO completed a single portfolio calculation in 10.90 seconds compared to 25.73 seconds for Zipline and 20.80 seconds for Rqalpha, representing a 136% and 91% speedup respectively. Under parallel execution with 24 processors, 121 parameter combinations were optimized simultaneously with near-linear efficiency (94% of theoretical linear scaling).

The portfolio constructed using the HPPO platform achieved a Sharpe ratio of 1.52 and an annualized return of 21.3% with maximum drawdown of -18.4%, outperforming standard Markowitz Mean-Variance Optimization (Sharpe 1.15). The information ratio of 0.89 indicates consistent active return generation relative to the benchmark.

Table 3: Computational Performance Comparison of Portfolio Platforms

Platform	Time per Calculation	Speedup vs. Zipline
Zipline	25.73s	Baseline
Rqalpha	20.80s	24% faster
HPPO [7]	10.90s	136% faster

VII. SYNTHESIS, CROSS-CUTTING THEMES, AND RESEARCH GAPS

7.1 The Case for Methodological Integration

A recurring observation across the surveyed literature is that technical price prediction and fundamental valuation have been pursued largely in isolation. Studies in the Transformer and LSTM literature rarely incorporate DCF-derived

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 7, March 2026
Impact Factor: 8.2

intrinsic value signals as model inputs, while fundamental analysis automation papers (e.g., [3]) do not typically leverage high-frequency price prediction as a complementary signal. Preliminary evidence from integrated system evaluations suggests that combining both streams yields portfolio performance improvements beyond what either stream achieves independently. The integrated system reported by the authors of the implementation work underlying this survey achieved a Sharpe ratio of 1.68 and annualized return of 23.7%, compared to 1.34 and 1.28 for standalone fundamental and technical components, respectively.

7.2 Identified Research Gaps

Based on systematic review of the selected literature, the following gaps are identified as priority directions for future research:

Integration gap: The majority of studies address either price prediction or portfolio construction, with few providing end-to-end frameworks that unify technical signals, fundamental valuation, and real-time execution within a single evaluated system.

Cross-market validation: Evaluations are predominantly conducted on U.S. equity indices. Generalization to emerging markets, different regulatory environments, and international exchanges remains underexplored.

Reproducibility: Many deep learning studies do not fully specify hyperparameter configurations, random seeds, or preprocessing pipelines, impeding replication and fair comparison.

Data quality and survivorship bias: Benchmark datasets typically exclude delisted securities, potentially overestimating historical backtest performance.

Real-time processing under constraints: Relatively few papers explicitly address the latency, throughput, and fault-tolerance requirements of production-grade real-time systems, a gap that the HPPO work [7] begins to address but does not fully resolve.

Model interpretability: Attention-based models offer limited financial interpretability relative to classical factor models. Developing attribution methods tailored to financial time series remains an open problem.

7.3 Methodological Trends and Convergence

The surveyed literature reflects a broader convergence toward attention mechanisms as the dominant architectural choice for sequential financial modeling, with Time2Vec and ProbSparse attention representing the current frontier. Simultaneously, reinforcement learning and imitation learning are gaining traction as principled alternatives to rule-based fundamental analysis. The parallel computing paradigm represented by HPPO is likely to become increasingly important as portfolio universes grow and regulatory pressure for real-time risk management intensifies.

An emerging methodological question concerns the optimal integration point between deep learning predictions and downstream portfolio optimization: whether to pass probabilistic forecasts as soft constraints, use predicted returns as inputs to mean-variance solvers, or train end-to-end systems that jointly optimize prediction and allocation objectives remains an active area of investigation.

VIII. CONCLUSION

This survey has reviewed the state of machine learning research in stock valuation and portfolio optimization across four principal domains: traditional DCF-based methods, deep learning price prediction architectures, imitation learning for automated fundamental analysis, and high-performance computing platforms for portfolio construction. The collective evidence from the reviewed studies indicates that machine learning methods have demonstrably advanced beyond classical benchmarks in all four areas, with Transformer-based models, imitation learning-augmented valuation, and distributed computing platforms each offering substantial empirical improvements over prior baselines.

Notwithstanding these advances, the field continues to be characterized by fragmentation between sub-disciplines, limited cross-market generalization, and reproducibility challenges. The most promising research trajectory identified in this survey is the systematic integration of technical and fundamental machine learning signals within

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 7, March 2026
Impact Factor: 8.2

computationally scalable frameworks. As data availability expands and model architectures continue to mature, the convergence of these streams is likely to yield both theoretical insights and practical systems of significant value to investment professionals and researchers alike.

REFERENCES

[1] J. Y. Lee and S. J. Yoo, "Stock Price Prediction Using Transformer and Time2Vec," in 2025 International Conference on Artificial Intelligence in Information and Communication (ICAIIC), IEEE, 2025, pp. 0687–0689, doi: 10.1109/ICAIIC64266.2025.10920805.

[2] R. Yulistiani and F. I. Kurniadi, "Stock Price Prediction with the Informer Model," in 2024 International Conference on Information Management and Technology (ICIMTech), Jakarta, Indonesia, IEEE, 2024, pp. 49–53, doi: 10.1109/ICIMTECH63123.2024.10780913.

[3] Y.-L. Peng and W.-P. Lee, "Valuation of Stocks by Integrating Discounted Cash Flow With Imitation Learning and Guided Policy," IEEE Transactions on Automation Science and Engineering, vol. 21, no. 2, pp. 1781–1792, April 2024, doi: 10.1109/TASE.2023.3246590.

[4] B. Math, O. Bevinakatti, S. Hosamani, and P. Patil, "Stock Market Prediction Using Machine Learning," in 2024 5th International Conference for Emerging Technology (INCET), Karnataka, India, IEEE, 2024, pp. 1–6, doi: 10.1109/INCET61516.2024.10593215.

[5] V. Goyal and R. R. Raj, "Using LSTM, DNN, and ARIMA approaches to predict the trends in the Stock Market," in 2022 8th International Conference on Advanced Computing and Communication Systems (ICACCS), Coimbatore, India, IEEE, 2022, pp. 918–922, doi: 10.1109/ICACCS54159.2022.9785150.

[6] J. Huang, "A Stock Price Prediction Method Based on Price Trend Curves," in 2024 17th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI), IEEE, 2024, pp. 1–6, doi: 10.1109/CISP-BMEI64163.2024.10906264.

[7] Y. Chen, Z. Lu, and X. Yang, "The Design and Implementation of a High-performance Portfolio Optimization Platform," in 2020 IEEE 23rd International Conference on Computational Science and Engineering (CSE), IEEE, 2020, pp. 1–7, doi: 10.1109/CSE50738.2020.00008.

[8] O. A. Malafeyev, A. V. Shuvaev, K. V. Yupatova, A. F. Dolgopolova, I. V. Zaitseva, and T. G. Gurnovich, "Social Network Valuation Using the Discounting Cash Flow Method," in 2021 1st International Conference on Technology Enhanced Learning in Higher Education (TELE), IEEE, 2021, pp. 139–142, doi: 10.1109/TELE52840.2021.9482554.

[9] M. Bansal, A. Goyal, and A. Choudhary, "Stock Market Prediction with High Accuracy using Machine Learning Techniques," Procedia Computer Science, vol. 210, pp. 351–358, 2022, doi: 10.1016/j.procs.2022.12.028.

[10] L. N. Mintarya, J. N. M. Halim, C. Angie, S. Achmad, and A. Kurniawan, "Machine learning approaches in stock market prediction: A systematic literature review," Procedia Computer Science, vol. 216, pp. 96–102, 2023, doi: 10.1016/j.procs.2022.12.115