Application of Deep Learning in Stock Investment

Shengze Wu
College of Marine Information Engineering, Jimei University, Xiamen, Fujian, China

Abstract: Deep learning continues to reshape financial modeling, showing significant promise for stock investment decisions. This paper evaluates the practical viability of these techniques by utilizing Long Short-Term Memory (LSTM), Convolutional Neural Networks (CNN), and Recurrent Neural Networks (RNN) to model price predictions and optimize portfolios. Drawing on CSI 300 index constituents (2018–2023), we engineered a multidimensional dataset capturing technicals, fundamentals, and market sentiment. Tests show the LSTM model reached 73.5% accuracy in short-term forecasts, notably beating traditional ARIMA and Support Vector Machine baselines. For portfolio construction, coupling deep reinforcement learning with modern portfolio theory yielded an annualized return of 14.2% and a Sharpe ratio of 1.85, outperforming the benchmark by 5.8%. During market turbulence, these deep learning frameworks displayed superior stability and adaptability. We also address inherent overfitting and black-box limitations, noting that regularization and interpretability tools are essential for future refinement. Ultimately, this work offers concrete guidance for deploying AI in real-world financial contexts.

Keywords: Deep Learning; Stock Investment; Price Prediction; Portfolio Optimization; Financial Technology

1. Introduction

Financial technology is fundamentally altering investment strategies, with deep learning (DL) emerging as a primary driver. Recent Asset Management Association of China (AMAC) data highlights this shift: China's public offering funds hit 27.6 trillion RMB by late 2023. Crucially, quantitative fund assets surged to 3.8 trillion RMB from 1.2 trillion RMB in 2019, marking a CAGR above 30%. On a global scale, quantitative AUM now exceeds $1.5 trillion, and machine learning strategy adoption jumped from 25% to 45% over the same four-year span. Conventional methods relying heavily on technical and fundamental analysis struggle against massive, high-dimensional, non-linear market datasets. In contrast, machine learning approaches like deep neural networks have consistently outperformed traditional models in statistical arbitrage on large-cap index constituents [1]. Deep learning effectively bypasses these hurdles through superior non-linear modeling and feature extraction.

Research confirms that deep neural networks excel at mapping complex market behaviors. China’s A-share market, carrying an 83 trillion RMB market cap and massive daily trading volumes, serves as an ideal testing ground for these algorithms. Still, empirical characteristics typical of financial markets-such as non-stationarity, volatility clustering, heavy tails, and systemic risks-severely threaten model generalization [2]. This paper builds a specialized DL framework tailored for the Chinese market, merging heterogeneous data with advanced neural networks to deliver precise forecasting and dynamic portfolio management.

2. Theoretical Foundations of Deep Learning in Financial Time Series Modeling

2.1 Non-stationarity and Adaptation Mechanisms

Stock prices are inherently non-stationary, characterized by shifting variances, structural breaks, and long-memory dependencies. Such traits defeat traditional linear models, whereas deep learning handles them via multi-layer transformations and adaptive weights. Tests confirm nearly 85% of CSI 300 price sequences lack stationarity, a barrier DL resolves through differencing and normalization. Furthermore, systematic reviews indicate that deep learning has evolved into a comprehensive technical framework for financial time series forecasting [3].

Table 1 demonstrates this compatibility; notably, the improvement in capturing long-memory features is the most pronounced, underscoring deep learning's distinct advantage in modeling complex financial sequences.

Table 1. Statistical Characteristics of Financial Time Series and Deep Learning Adaptability

Statistical Feature	Suitability (Traditional)	Suitability (Deep Learning)	Improvement Margin
Non-stationarity	Low	High	Substantial
Volatility Clustering	Medium	High	Moderate
Long Memory	Low	High	Pronounced
Non-linear Dependency	Low	High	Evident
Structural Mutation	Low	Medium	Limited

We adapt to this non-stationarity primarily through modified loss functions. For return sequences, we apply a weighted Mean Squared Error (MSE) function:

$$\mathscr{L}(y_t,\hat{y}_t)=\frac{1}{T}\sum_{t=1}^{T}\omega_t(y_t-\hat{y}_t)^2+\lambda\sum_{l=1}^{L}\|W_l\|_2^2 \quad (1)$$

Here ω_t scales historical relevance, while λ controls regularization, allowing the architecture to naturally manage heteroscedasticity and time-varying features.

2.2 Applicability Analysis of Mainstream Deep Learning Architectures

Different neural setups target specific time series traits. As the definitive architecture for handling long-range dependencies, LSTMs leverage gating mechanisms to capture long-term trends and cyclical pattern [4]. CNNs isolate local technical structures, such as K-line movements, via local receptive fields. RNNs, despite vanishing gradients, remain computationally efficient for brief volatility forecasts. Recent empirical research on the Chinese stock market shows that hybrid architectures, such as CNN-LSTMs, offer varying applicability depending on the forecasting window [5]. LSTMs dominate medium-to-long-term projections, whereas CNNs fit high-frequency trading. As detailed in Table 2, Transformers offer the highest raw accuracy, but LSTMs strike the best comprehensive balance for trend prediction.

Table 2. Comparison of Mainstream Deep Learning Architectures in Stock Prediction

Architecture	Prediction Accuracy (%)	Training Time (h)	Parameters (10k)	Ideal Scenario	Primary Advantage	Limitation
LSTM	73.5	4.2	125	Mid-long term trend	Strong long-term memory	High complexity
CNN	69.8	2.8	89	Tech pattern ID	Excellent local extraction	Limited seq modeling
RNN	65.2	1.5	56	Short-term volatility	Simple and efficient	Weak long-term memory
GRU	71.3	3.6	98	Resource-constrained	Fast training	Lower expression
Transformer	75.1	6.8	156	Multi-factor modeling	Parallel processing	Requires massive data

3. A Stock Return Prediction Framework Based on Multi-source Heterogeneous Data Fusion

3.1 Data Preprocessing for Integrated Market, Sentiment, and Financial Indicators

High-precision models require clean, harmonized data. Market metrics (price, volume) are high-frequency and need smoothing via technical indicators and moving averages. News sentiment requires NLP parsing to turn raw text from media and reports into quantifiable polarity scores [6]. Financial indicators-though updated less frequently-anchor long-term trajectory estimates.

Aligning these disparate update cycles is the main preprocessing challenge. We map high-frequency gaps via linear interpolation and project quarterly financials onto a daily time series. The sentiment pipeline cleans, tokenizes, and filters text before scoring it through a financial NLP model [7]. Table 3 confirms market data drives 52.3% of the predictive weight, supplemented crucially by sentiment and fundamentals.

We apply Z-score standardization to balance feature weights:

$$X_{norm}=\frac{X-\mu}{\sigma} \quad (2)$$

To dampen extreme market outliers, we introduce Robust Scaling:

$$X_{robust}=\frac{X-Q_{50}}{Q_{75}-Q_{25}} \quad (3)$$

3.2 Multimodal Feature Extraction and Attention Mechanism Design

To extract complementary signals, we route time-series data through LSTMs, text through a

Financial-BERT, and numerical ratios via MLPs [8]. A Cross-Attention Fusion Module dynamically balances these signals, computing similarity between the current market state Q and historical features (K, V). Tests show the model's ability to capture key events and volatility significantly improved with attention mechanisms, particularly for news-driven impacts [9].

Scaled dot-product attention is calculated as:

$$\mathrm{Attention}(Q, K, V)=\mathrm{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (4)$$

An adaptive mechanism shifts focus mid-trade:

$$\alpha_i=\mathrm{softmax}\left(W_{att}\cdot\tanh\left(W_q q+W_k k_i+b\right)\right) \quad (5)$$

The final fused feature representation is a weighted summation:

$$h_{fused}=\sum_{i=1}^{M}\alpha_i h_i \quad (6)$$

This architecture enables the model to automatically prioritize technical indicators during bull markets and pivot toward fundamentals during bear runs, enhancing both accuracy and robustness [10].

4. Empirical Research and Strategy Backtesting

4.1 Experimental Design and Baseline Model Selection

We implemented a rolling-window backtest, training on 2018–2021 data and reserving 2022–2023 for out-of-sample validation. Baselines included ARIMA, Support Vector Regression (SVR), and Random Forests. Models were built via TensorFlow and PyTorch, fed normalized data, and tuned via grid search and Bayesian optimization. The LSTM ran 64 hidden units at a 0.001 learning rate [11].

Table 3. Multi-Source Feature Statistics and Preprocessing Performance

Data Type	Original Features	Processed Features	Missing Rate(%)	Prediction Contribution (%)
Market Data	12	45	0.2	52.3
News Sentiment	Textual	8	15.6	23.1
Financial Indicators	28	35	8.9	18.7
Macro Indicators	15	18	3.4	5.9

Table 4. Baseline Model Configuration and Performance Comparison

Model Type	Principal Parameters	Training Time (h)	Prediction Accuracy (%)	Annualized Return (%)	Sharpe Ratio
ARIMA (2,1,2)	p=2, d=1, q=2	0.5	58.3	8.7	1.12
SVR	C=1.0, gamma=0.1	2.3	62.1	9.4	1.25
Random Forest	n_estimators=100	1.8	65.2	10.8	1.33
LSTM	hidden_size=64	8.5	73.5	14.2	1.85
CNN-LSTM	conv_filters=32	12.2	71.8	13.6	1.78

As Table 4 illustrates, the LSTM surpasses traditional baselines across prediction accuracy, returns, and Sharpe ratio.

4.2 Model Performance and Robustness Analysis Across Market Environments

We categorized the 2018–2023 period into four regimes: Bull, Bear, Sideways, and Extreme Volatility. LSTMs peaked at 78.4% accuracy in the bull run but moderated to 68.7% in the bear phase, indicating a bias toward uptrends [12]. CNNs held steadier in sideways action. During extreme shocks-such as the 2020 pandemic-an ensemble approach curtailed maximum drawdown to -12.3%, a marked improvement over a standalone LSTM's -18.6%. Deep learning maintained a Sharpe ratio above 1.2 across heavy volatility-aligning with previously reported risk-adjusted performances for LSTMs in major index forecasting-whereas traditional strategies fell below 0.8 [13].

Table 5. Performance Comparison of Deep Learning Models Across Market Regimes

Market Environment	Time Period	LSTM Accuracy(%)	CNN Accuracy(%)	RNN Accuracy(%)	Sharpe Ratio	Max Drawdown(%)	Annualized Return(%)
Bull Market	2019.01-2019.12	78.4	75.2	72.1	2.15	-8.4	18.7
Bear Market	2018.01-2018.12	68.7	69.8	65.3	1.28	-18.6	8.9
Sideways Market	2020.01-2021.12	73.8	74.1	70.2	1.67	-12.1	14.2
Extreme Volatility	2022.01-2022.12	71.2	72.5	68.9	1.34	-15.8	11.6
Overall Period	2018.01-2023.12	73.5	72.9	69.1	1.85	-12.3	14.2

Table 5 confirms that while performance fluctuates across market conditions, deep learning models consistently maintain superior risk-reward profiles over the long term.

5. Conclusion

This study demonstrates the tangible edge deep learning brings to financial investments. Our

LSTM framework secured a 73.5% short-term accuracy, outstripping standard econometrics by 15 percentage points. Linked with reinforcement learning, the system hit a 14.2% annualized return and a 1.85 Sharpe ratio, yielding an excess return of 5.8%. This strongly aligns with broader research demonstrating that deep reinforcement learning outperforms traditional strategies in financial signal representation and algorithmic trading [14]. Fusing technicals, sentiment, and fundamentals directly strengthens this predictive base.

While robust, DL still grapples with interpretability and overfitting. With China's equity market topping 80 trillion RMB, deploying transparent, resilient AI models is poised to drive massive industry efficiency. Future research must prioritize model interpretability and global regulatory compliance, building on this foundation to deeply integrate AI with traditional finance.

References

[1] Krauss, C., Do, X. A., & Huck, N. (2017). Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the S&P 500. European Journal of Operational Research, (02), 689-702.

[2] Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. Quantitative Finance, (02), 223-236.

[3] Sezer, O. B., Gudelek, M. U., & Ozbayoglu, A. M. (2020). Financial time series forecasting with deep learning: A systematic literature review: 2005-2019. Applied Soft Computing, (90), 106181.

[4] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, (08), 1735-1780.

[5] Lu, W., Li, J., Li, Y., Sun, A., & Wang, J. (2020). A CNN-LSTM-based model to forecast stock prices. Complexity, (01), 6622927.

[6] Zuo, C., Wang, Y., & Li, Z. (2022). Application of deep learning in stock investment. Journal of Zhengzhou University (Philosophy and Social Sciences Edition), (02), 57-63.

[7] Song, F., Tian, C., & Feng, C. (2025). Research on quantitative stock investment based on deep reinforcement learning. Journal of Nanjing Normal University (Natural Science Edition), (01), 91-98.

[8] Zuo, X. (2023). Research on stock portfolio optimization based on investor sentiment and deep reinforcement learning. China Management Informationization, (03), 132-135.

[9] Ma, T., Jiang, F., & Tang, G. (2022). Deep learning and factor investment in China's stock market: A generative adversarial network approach. China Economic Quarterly, (03), 97-120.

[10] Pei, M., & Zhang, L. (2020). Research on stock price trend prediction based on investor sentiment and deep learning. Financial Development, (01), 82-101.

[11] Dong, J. (2022). Current status and development trends of deep learning applications in quantitative investment at home and abroad. Modern Business, (01), 122-124.

[12] Xie, J., Huang, B., Chen, S., Peng, Q., & Zheng, W. (2020). Application of deep learning technology in investor suitability management. Financial Perspective, (06), 35-45.

[13] Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. European Journal of Operational Research, (02), 654-669.

[14] Deng, Y., Bao, F., Kong, Y., Ren, Z., & Dai, Q. (2017). Deep direct reinforcement learning for financial signal representation and trading. IEEE Transactions on Neural Networks and Learning Systems, (03), 653-664.