Stock Price Prediction Based on Standardized Price–Volume Charts and a Convolutional Neural Network

Tao Lin[1], Zhuming Chen[2, *], Ningning An[1], Yuanwen Chen[1]

[1] Research Centre and Development, GF Securities, Guangzhou, 510627, China
[2] Research Center for Digital Assets and Digital Finance, School of Accounting, Nanfang College, Guangzhou, 510970, China
[*]Email: chenzhm@mail.sysu.edu.cn (Corresponding Author)

Abstract
Forecasting stock price movements from price–volume information remains challenging yet widely studied. Using daily A-share data from 2005 to 2025, we construct standardized price–volume charts and adopt a computer-vision approach to prediction. Specifically, we transform each stock's price–volume history into a fixed-format chart representation and train a Convolutional Neural Network (CNN) to map chart patterns to subsequent price movements. We evaluate the model using 2005–2019 as in-sample data and 2020–2025 as an out-of-sample test period. Empirical results demonstrate that the CNN effectively extracts morphological features to predict future trends. In the out-of-sample period (2020–2025), the long portfolio constructed based on the model's predictions achieves an annualized return of 18.20% in the broad A-share market, generating an annualized excess return of 13.86% relative to the CSI All Share Index. The approach proves particularly effective in the volatile ChiNext sector, yielding an annualized return of 23.54% (13.22% excess return) and a mean Rank IC of 6.2%. Furthermore, the strategy remains profitable after accounting for transaction costs. Overall, our findings indicate that encoding price–volume data as standardized charts can be a useful complement to sequence-based modeling in equity prediction tasks.

Keywords: Deep learning; Convolutional neural network; standardized price–volume charts; trend classification; China A-share market; portfolio backtesting

Received: March 20, 2025
Revised: May 01, 2025
Accepted: June 03, 2025
Available Online: June 30, 2025

See: https://inatgi.in/index.php/jbda/index for more information. DOI:10.63646/j.jbda.2025.030201

Stock Price Prediction Based on Standardized Price–Volume Charts and a Convolutional Neural Network

1. Introduction

Despite the dominance of algorithmic trading, a significant portion of market liquidity—particularly in emerging markets like China's A-shares—is driven by decisions made through the visual inspection of price charts. While human traders synthesize complex patterns from candlesticks, moving averages, and volume bars instantly (a process known as "chart morphology" analysis), most current deep learning literature treats stock prediction strictly as a numeric sequence modeling problem. This creates a critical "Cognitive Gap": state-of-the-art sequence models (RNNs, Transformers) may capture temporal correlations, but they fail to capture the spatial and topological structures that drive technical analysis and market sentiment.

The urgency of addressing this gap is underscored by the limitations of sequence-based models in volatile markets. Recent studies indicate that while models like LSTM and Transformers excel at capturing long-term dependencies, they struggle to represent the "shape" information—how price and volume co-move to form recognizable patterns (e.g., "Golden Cross" or "Divergence")—that dictates short-term trend reversals. Jiang et al. (2023) recently argued for "re-imaging" price trends, suggesting that visual representations can capture non-linear interactions between price and volume that 1D time-series models overlook. However, existing image-based approaches often rely on simple K-line screenshots, lacking the rich, multi-dimensional data (Moving Averages, MACD, Turnover) required for professional-grade analysis.

To bridge this gap, this paper proposes a novel framework that shifts the paradigm from "sequence modeling on raw numbers" to "pattern recognition on standardized charts." We introduce a strictly specified Standardized Price–Volume Chart encoding scheme that converts multi-dimensional market data into a fixed-format visual tensor. This standardization is the critical innovation: it ensures that specific pixel regions always correspond to specific information types (Price, Turnover, MACD), allowing a Convolutional Neural Network (CNN) to learn consistent, robust morphological features across thousands of stocks and two decades of data.

Our contributions are threefold:

1. Methodological Innovation: We propose a reproducible "Standardized Chart" representation that embeds high-dimensional technical indicators (MA lines, MACD, Volume) into a single 2D image, overcoming the information loss inherent in simple K-line plotting.
2. Theoretical Shift: We demonstrate that stock prediction can be effectively modeled as a Computer Vision (CV) problem, utilizing CNNs to extract spatial features that mimic the cognitive process of technical analysts.
3. Robust Empirical Evidence: Unlike studies focused on short windows, we evaluate our model over a 20-year horizon (2005–2025) in the China A-share market. Out-of-sample testing (2020–2025) confirms that our visual-learning approach delivers significant excess returns, particularly in retail-dominated sectors like ChiNext, validating the hypothesis that visual chart patterns contain alpha uncaptured by traditional sequence models.

2. Related Work

According to the Efficient Market Hypothesis (EMH) proposed by Fama et al. (1970), in an efficient market stock prices incorporate all relevant information; as a result, price changes are unpredictable. However, Abu-Mostafa et al. (1996) argue that stock markets exhibit price trends and correlations between fundamental events and economic data. Behavioral finance likewise suggests that markets are not perfectly efficient, implying that stock prices may be predictable to a limited extent (Zhong et al., 2017). Park et al. (2007) and Nguyen et al. (2015) classify stock analysis into fundamental analysis and technical analysis.

Hu et al. (2015) posit that fundamental analysis is grounded in three levels: (1) Macroeconomic analysis, which examines the effects of macro conditions on future corporate profitability using indicators such as GDP, CPI, and PMI; (2) Industry analysis, which evaluates firm value based on industry status and prospects; and (3) Company analysis, which assesses a firm's operating and financial condition to estimate intrinsic value and benchmark it against peers. Although fundamental analysis can offer insight into future trends, fundamental data are often low-frequency and released with a lag relative to continuously traded markets (e.g., China's GDP is quarterly, CPI/PMI monthly, and financial statements quarterly). By contrast, technical analysis is primarily based on price–volume data and can leverage daily or even real-time information to extract signals that track market dynamics more closely.

Early studies on stock prediction using price–volume data focused on non–deep learning methods. These approaches typically start from raw price–volume time series and employ models such as ARMA, ARIMA, ARCH, GARCH, SVM, KNN, Random Forest, and XGBoost (Tyssedal et al., 1988; Pai et al., 2005; Guresen et al., 2011; Ballings et al., 2015; Dey et al., 2016; Imandoust et al., 2013; Zhong et al., 2017; Shah et al., 2019).

In recent years, rapid advances in deep learning have created new opportunities for stock prediction from price–volume data. Di Persio and Honchar (2017) apply RNN-based models, including LSTM and GRU, to forecast Google's stock price. Yang et al. (2017) employ an ensemble of multilayer feedforward neural networks for the Chinese market. Zhao and Xue (2021) propose a hybrid LSTM-CNN-CBAM model that combines sequence modeling with convolutional feature extraction and enhances prediction via a CBAM attention module. Zhou and He (2023) predict sector-index trends using a PCA-based representation and an attention-optimized LSTM (PCA-Attention-LSTM).

Beyond RNN-based architecture, the research focus has shifted significantly toward Transformer models and self-attention mechanisms in the post-2020 era. Limitations in the parallelization and long-term memory of LSTMs have led researchers to adopt architectures such as the Temporal Fusion Transformer (TFT) (Lim et al., 2021) and Informer (Zhou et al., 2021), which excel at capturing long-range dependencies in volatile markets. For instance, recent studies have demonstrated that multi-head attention mechanisms can effectively isolate critical market events from noise, offering superior performance over standard recursive models by enhancing locality and breaking memory bottlenecks (Li et al., 2019; Vaswani et al., 2017). However, despite their ability to model temporal dynamics, these sequence-based Transformers often treat price and volume as purely numerical inputs, neglecting the topological and morphological structures inherent in technical charting.

In summary, stock predictions remain a challenging task, both academically and in practice. Ongoing progress in AI has strengthened the ability to learn from price–volume information, but much of the literature still concentrates on modeling raw sequences or relying on relatively simple K-line technical analysis, which can be limiting. Motivated by these gaps, we explore a chart-based approach that standardizes price–volume information into a visual representation and applies CNNs to recognize price-trend morphology.

3. Standardized Price-Volume Charts

To construct the standardized price-volume charts, we use daily forward-adjusted data for all stocks in the Shanghai and Shenzhen A-share markets from January 1, 2005, to December 31, 2025, including the opening, low, and closing prices, trading volume, and turnover amount. All data is obtained from the TinySoft database.

Based on the closing price, we compute moving averages (MA) with window lengths of 5, 10, 20, 60, 120, and 250 trading days. We also compute the Moving Average Convergence Divergence (MACD) indicator from the closing price, using 12 and 26 days for the fast and slow lines, respectively, and 9 days for the signal line (DEA). The basic price–volume fields and the derived indicators together yield 15 series used to generate the standardized charts.

Table 1: Basic and Derived Data for Standardized Price-Volume Charts

Variable Name	Meaning
Open	Opening Price
High	Highest Price
Low	Lowest Price
Close	Closing Price
Volume	Trading Volume
Amount	Turnover Amount
MA5	5-day Avg Price
MA10	10-day Avg Price
MA20	20-day Avg Price
MA60	60-day Avg Price
MA120	120-day Avg Price
MA250	250-day Avg Price
MACD	MACD Value
DIF	Difference Line / Fast Line
DEA	Signal Line / Slow Line

We refer to the resulting images as “standardized” because both the rendering format and the data layout follow a single specification, enabling the convolutional neural network to learn comparable price–volume morphologies across stocks and time. As shown in Figure 1, each standardized chart is rendered at 200 × 200 pixels with a gray background and summarizes the 15 series over a 20-day window, arranged into three stacked regions:

(a) Upper region (120 pixels). We plot the candlestick (K-line) chart together with the six MA lines. Following the A-share convention, red and green indicate price increases and

See: https://inatgi.in/index.php/jbda/index for more information. DOI:10.63646/j.jbda.2025.030201

decreases, respectively, and the MA lines are drawn in fixed colors to support visual comparability.

(b) Middle region (40 pixels). We display turnover using bars whose colors match the day's candlestick direction. Although the price series are forward adjusted, we use turnover amount (rather than trading volume) in this panel to reduce distortions introduced by changes in share capital (e.g., stock splits), which mechanically affect raw volume.

(c) Lower region (40 pixels). We visualize MACD-related information: MACD is shown as a bar chart, while DIF and DEA are plotted as line charts using fixed colors.

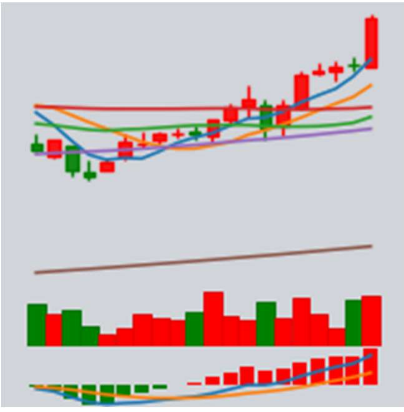

Figure 1: Standardized Price-Volume Chart

4. Efficient Stock Price Prediction Convolutional Neural Network

The overall structure of the proposed Efficient Stock Price Prediction CNN is shown in Figure 2, with the standardized price-volume chart denoted as Xserving as the input. Digital images consist of raw pixel values ranging from 0 to 255 across three channels (Red, Green, and Blue). Accordingly, each standardized price–volume chart is represented as a tensor of size $3 \times 200 \times 200$, where 3 corresponds to the RGB channels and 200×200denotes the image resolution.

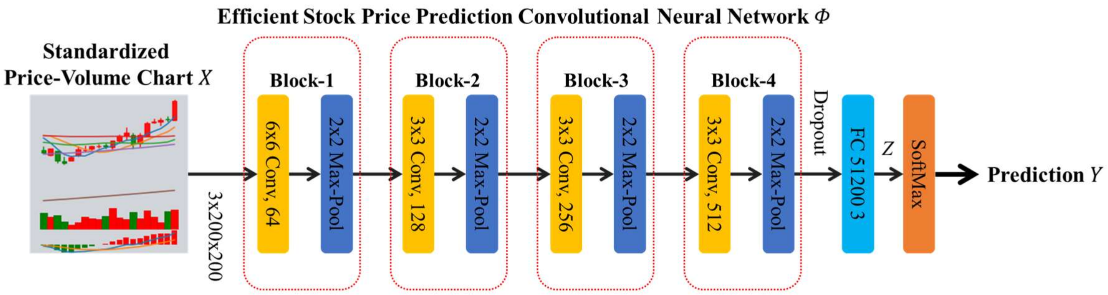

Figure 2: Efficient Stock Price Prediction Convolutional Neural Network

The backbone of the model consists of four convolutional blocks. Each block includes a Convolutional layer (Conv), Batch Normalization (BatchNorm), an LReLU activation, and a Max Pooling layer. The convolutional layer performs feature extraction through local convolution operations (i.e., elementwise multiplication between the kernel and input patches followed by summation). Batch normalization is then applied to normalize

See: https://inatgi.in/index.php/jbda/index for more information. DOI:10.63646/j.jbda.2025.030201

intermediate activations, which typically improves optimization stability and training speed and can help suppress overfitting:

$$y = \frac{x - E(x)}{\sqrt{Var(x) + \varepsilon}} * \alpha + \beta \tag{1}$$

where $x$is the input, $\mathbb{E}(x)$is the mean, $\mathrm{Var}(x)$is the variance, $\varepsilon$is a small constant to avoid division by zero, and $\alpha$and $\beta$are learnable parameters. After normalization, we apply the Leaky Rectified Linear Unit (LReLU):

$$LReLU(x) = \begin{cases} x, & x > 0 \\ slope * x, & x \leq 0 \end{cases} \tag{2}$$

The overall structure of the proposed Efficient Stock Price Prediction CNN is shown in **Figure 2**, with the standardized price-volume chart denoted as X serving as the input. Typically, digital images consist of raw pixel values ranging from 0 to 255, composed of Red, Green, and Blue (RGB) channels. Therefore, the matrix dimension of the standardized price-volume chart is $3 \times 200 \times 200$, where 3 represents the RGB channels, and 200×200 represents the pixel dimensions.

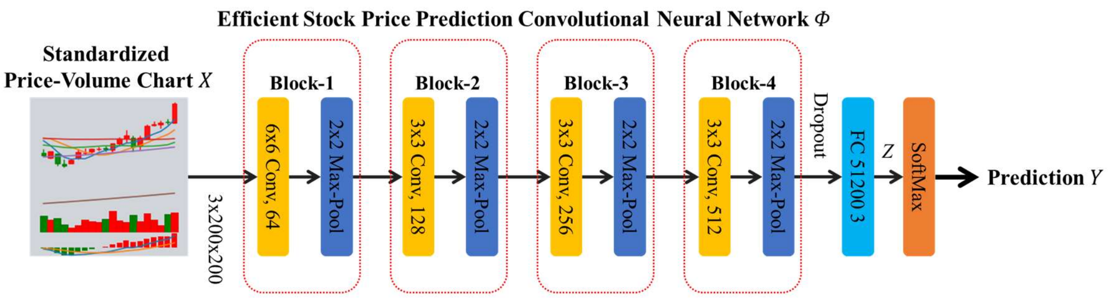

Figure 2: Efficient Stock Price Prediction Convolutional Neural Network

The backbone of the model consists of 4 convolutional blocks (Block). Each block comprises a Convolutional Layer (Conv), Batch Normalization (BatchNorm), LReLU activation function, and Max Pooling layer. The convolutional layer is the core structure, extracting features via convolution operations (multiplication and addition of the convolution kernel with input data). Batch Normalization (Ioffe, 2015) is then used to map the output of the convolution kernel to a distribution with a mean of 0 and variance of 1, accelerating training and suppressing overfitting:

$$y = \frac{x - E(x)}{\sqrt{Var(x) + \varepsilon}} * \alpha + \beta \tag{1}$$

where x is the input, $E(x)$ is the mean, $Var(x)$ is the variance, ε is a small value to prevent division by zero, and α and β are learnable parameters. After normalization, we apply the Leaky Rectified Linear Unit (LReLU):

See: https://inatgi.in/index.php/jbda/index for more information. DOI:10.63646/j.jbda.2025.030201

$$LReLU(x) = \begin{cases} x, & x > 0 \\ slope * x, & x \leq 0 \end{cases} \tag{2}$$

where slope is a small constant (typically 0.0176). Unlike standard ReLU, which outputs zero for negative inputs (potentially leading to "dead neurons"), LReLU allows a small negative slope so that gradients can still propagate. Finally, Max Pooling expands the receptive field and reduces spatial resolution, which can reduce overfitting.

In Block-1, the network extracts a $64 \times 97 \times 97$ feature map. In Blocks 2–4, the representation is progressively refined, and a $512 \times 10 \times 10$ feature map is obtained before being mapped to the output through a fully connected (FC) layer. The output layer has three nodes corresponding to three future price-trend classes: Up (positive excess return), Flat (no excess return), and Down (negative excess return).

The network models the mapping between the input chart X and the output Y. Let Φ denote the model parameters and let $Z = [z_1, z_2, z_3]^{\mathrm{T}}$ be the FC-layer output. We convert logits into class probabilities using Softmax:

$$\hat{Y} = [\hat{y}_1 \quad \hat{y}_2 \quad \hat{y}_3]^T = \left[\frac{e^{z_1}}{\sum_{i=1}^{3} e^{z_i}} \quad \frac{e^{z_2}}{\sum_{i=1}^{3} e^{z_i}} \quad \frac{e^{z_3}}{\sum_{i=1}^{3} e^{z_i}}\right]^T \tag{3}$$

For optimization, we adopt the Cross-Entropy loss:

$$E(\Phi) = -\sum_{n=1}^{N} \sum_{k=1}^{K} [y_{nk} \log(\hat{y}_{nk}) + (1 - y_{nk}) \log(1 - \hat{y}_{nk})] \tag{4}$$

where y_{nk} denotes the actual class indicator and $\hat{y}_{nk}$ is the predicted probability. Parameters are optimized using the Adam optimizer.

The model is trained on all A-share stocks from January 2005 to December 2016. We use a validation set from February 2017 to November 2019 to determine the early-stopping point. To prevent information leakage, we impose a buffer period of one calendar month after the training phase. Because the prediction target is the return over the subsequent 20 trading days, this buffer (longer than the horizon) ensures that training labels do not overlap with validation inputs. For image preprocessing, we compute the mean and standard deviation of the three RGB channels using only the training set and then apply Z-score normalization to all inputs. Out-of-sample evaluation is conducted from January 2020 to December 2025, and we exclude stocks with fewer than 250 trading days since listing across all splits.

5. Empirical Analysis

5.1 CNN Feature Visualization

Following training, we use the standardized price–volume chart in Figure 1 as a representative input and visualize the model's internal representations. Feature maps are randomly sampled from the outputs of the four convolutional blocks and summarized in Figure 3. The results suggest a clear hierarchy of representations: Convolutional Blocks 1–

See: https://inatgi.in/index.php/jbda/index for more information. DOI:10.63646/j.jbda.2025.030201

2 behave as low-level feature extractors, capturing global chart elements such as candlesticks, moving averages, turnover, and MACD, whereas Blocks 3–4 exhibit more localized and differentiated attention across the chart regions (e.g., some maps emphasize candlestick/MA patterns, others emphasize turnover or MACD, while a subset remains globally responsive). Overall, these visualizations are consistent with the intended role of the CNN as a morphology-based pattern learner on the standardized charts.

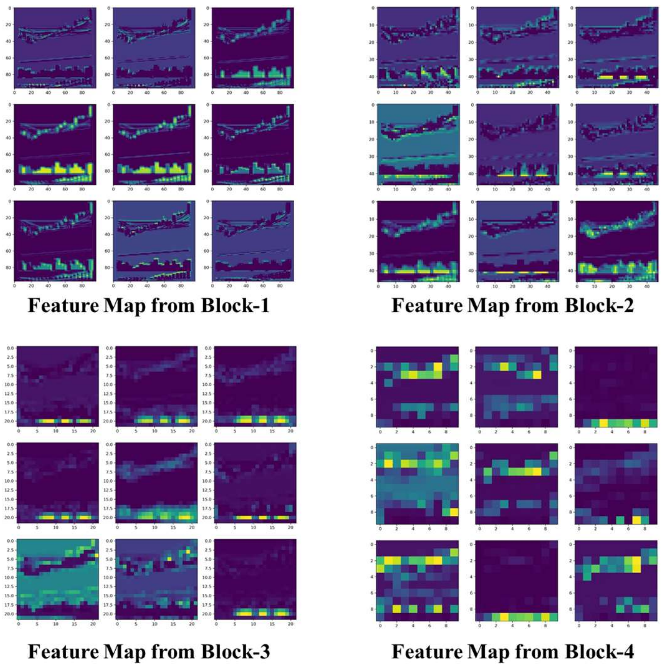

Figure 3: Visualization of CNN Features

5.2 Sub-sector Empirical Analysis

To evaluate whether the model outputs translate into implementable trading signals, we form portfolios using the predicted probability of an upward movement. Stocks are sorted cross-sectionally by the model's upward probability and grouped into five quintiles (Q1–Q5) in descending order; Q1 is treated as the long portfolio. Portfolios are equally weighted and rebalanced every 20 trading days. To improve realism, we account for transaction frictions (slippage and market impact) and assume execution at the next day VWAP (T+1) based on the signal generated on day T. In addition, we report a robust comparison that applies a bilateral transaction fee of 0.3%. All results below are based on out-of-sample data from January 2020 to December 2025, and each segment is benchmarked against its corresponding sector index.

5.2.1 CSI All Share

In the full A-share universe, the Q1 long portfolio achieves an annualized return of 18.20%, compared with 4.34% for the CSI All Share Index, implying an annualized excess return of 13.86%. The Q1 portfolio's average turnover is 73.42%, and the annualized return remains 15.20% after applying the 0.3% bilateral fee. The cumulative-return stratification across quintiles is summarized in Figure 4, and the association between predicted upward probabilities and realized returns (measured by Spearman rank correlation / Rank IC) is

See: https://inatgi.in/index.php/jbda/index for more information. DOI:10.63646/j.jbda.2025.030201

summarized in Figure 5; the historical average Rank IC is 4.8%, with a cumulative profile that is predominantly positive.

Table 2: Performance Statistics of the Entire A-share Market

	Annualized Return	Max Drawdown	Annualized Volatility	Sharpe Ratio
Long(Q1)	18.20%	29.20%	20.00%	0.79
Long(Q1) (with 0.3% Fee)	15.20%	29.85%	20.00%	0.64
CSI All Share Index	4.34%	39.33%	19.55%	0.09
Q1 Excess Return (vs. CSI All Share Index)	13.86%	17.97%	16.85%	0.55

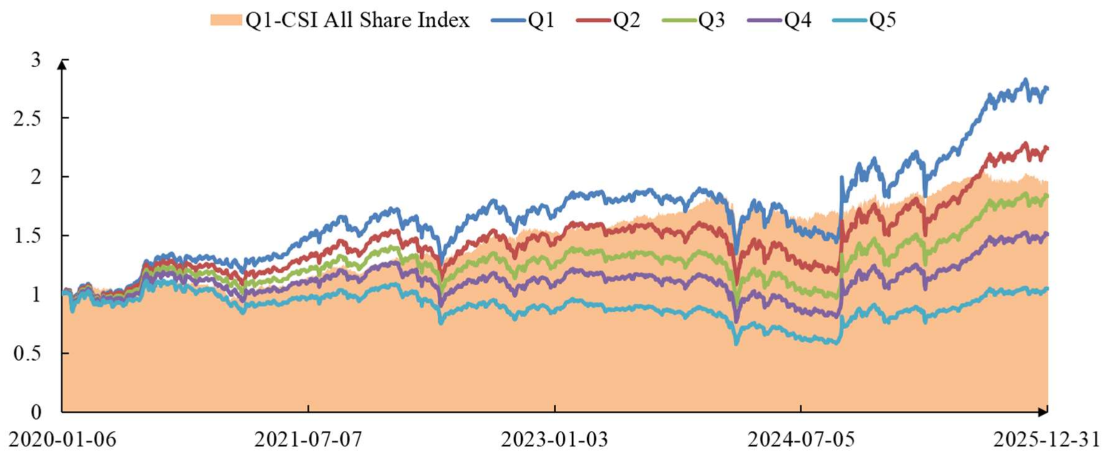

Figure 4: Cumulative Return of the Entire A-share Market

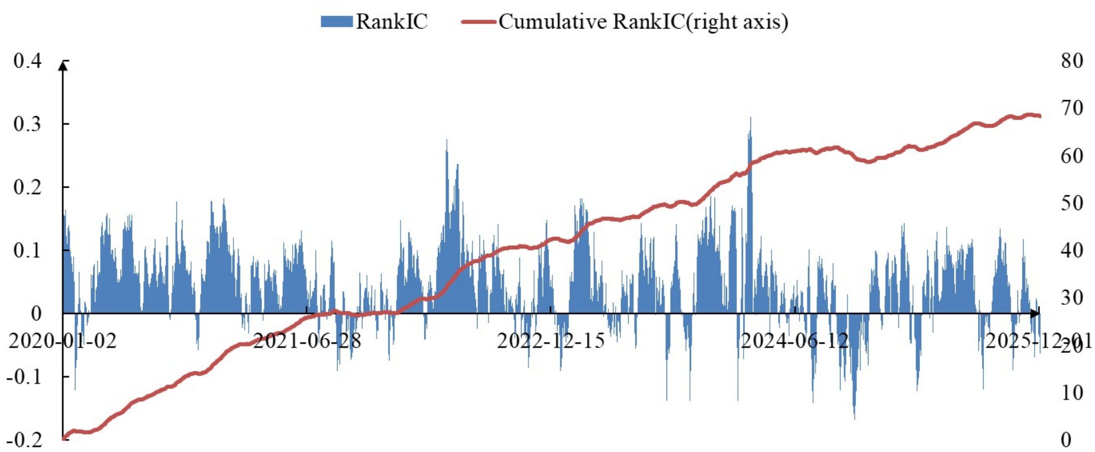

Figure 5: Cumulative Rank IC of the Entire A-share Market

5.2.2 CSI 300

See: https://inatgi.in/index.php/jbda/index for more information. DOI:10.63646/j.jbda.2025.030201

Within the CSI 300, the Q1 long portfolio attains an annualized return of 10.12%, versus 1.85% for the CSI 300 Index, corresponding to an excess return of 8.27%. The Q1 turnover averages 73.39%, and the annualized return is 7.33% after the 0.3% bilateral fee. Figures 6 and 7 summarize the cumulative performance and the Rank-IC relationship; the historical average Rank IC is 4.0%, and the cumulative Rank IC shows a broadly increasing tendency over time.

Table 3: Performance Statistics of the CSI 300

	Annualized Return	Max Drawdown	Annualized Volatility	Sharpe Ratio
Long(Q1)	10.12%	30.34%	16.77%	0.45
Long(Q1) (with 0.3% Fee)	7.33%	33.87%	16.77%	0.29
CSI 300 Index	1.85%	45.60%	18.49%	-0.04
Q1 Excess Return (vs. CSI 300 Index)	8.27%	7.75%	13.64%	0.32

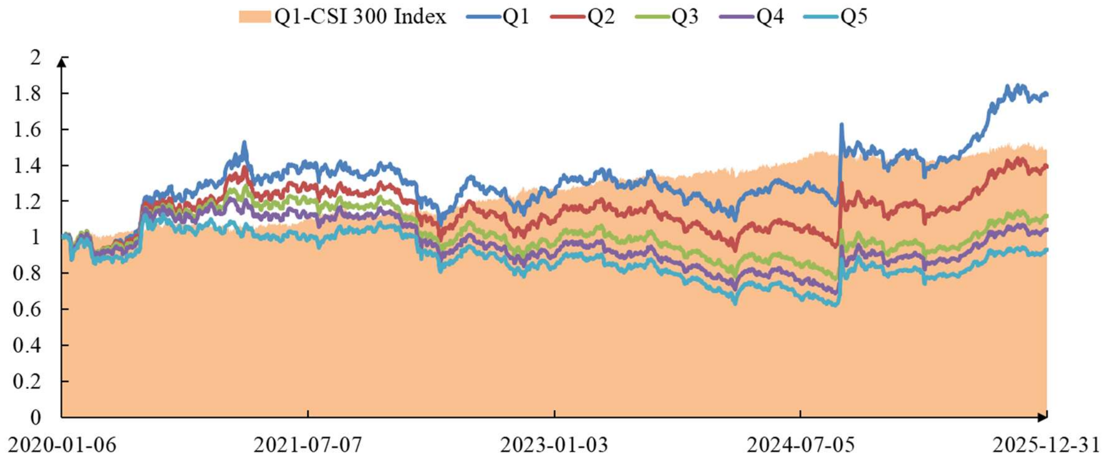

Figure 6: Cumulative Return of the CSI 300

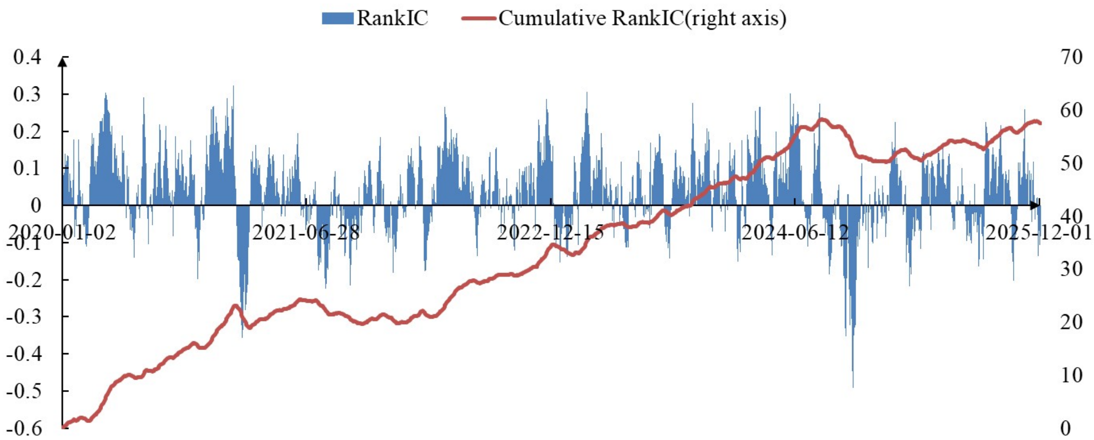

Figure 7: Cumulative Rank IC of the CSI 300

See: https://inatgi.in/index.php/jbda/index for more information. DOI:10.63646/j.jbda.2025.030201

5.2.3 CSI 500

For the CSI 500, Q1 records an annualized return of 12.47%, compared with 5.56% for the CSI 500 Index, yielding an excess return of 6.91%. The Q1 turnover is 73.33%, and the annualized return is 9.61% after applying the 0.3% bilateral fee. Figures 8 and 9 report the cumulative return stratification and the Rank-IC series; the historical average Rank IC is 3.0%, and the cumulative Rank IC exhibits a modest monotonic increase.

Table 4: Performance Statistics of the CSI 500

	Annualized Return	Max Drawdown	Annualized Volatility	Sharpe Ratio
Long(Q1)	12.47%	31.24%	18.06%	0.55
Long(Q1) (with 0.3% Fee)	9.61%	33.84%	18.06%	0.39
CSI 500 Index	5.56%	41.81%	21.06%	0.15
Q1 Excess Return (vs. CSI 500 Index)	6.91%	9.10%	15.21%	0.14

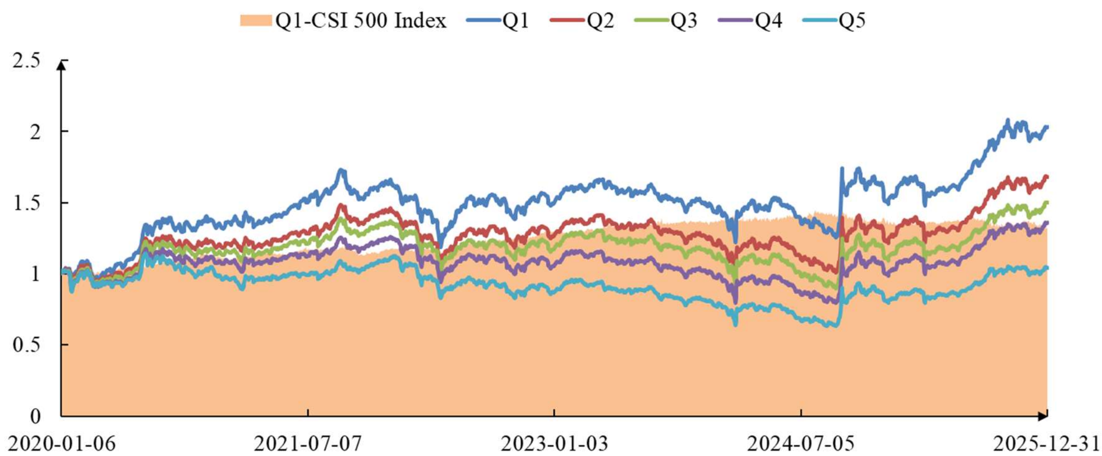

Figure 8: Cumulative Return of the CSI 500

See: https://inatgi.in/index.php/jbda/index for more information. DOI:10.63646/j.jbda.2025.030201

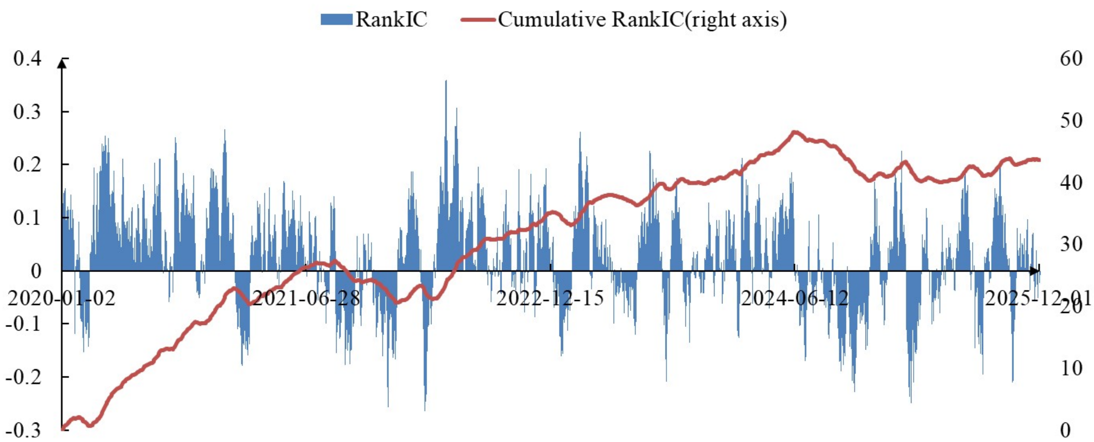

Figure 9: Cumulative Rank IC of the CSI 500

5.2.4 CSI 1000

For the CSI 1000, Q1 achieves an annualized return of 14.02%, while the CSI 1000 Index returns 4.84%, implying an excess return of 9.18%. The Q1 turnover averages 74.17%, and the annualized return is 11.10% after the 0.3% bilateral fee. Figures 10 and 11 summarize cumulative returns and the Rank-IC relationship; the historical average Rank IC is 4.0%, and the cumulative Rank IC shows a generally increasing trajectory.

Table 5: Performance Statistics of the CSI 1000

	Annualized Return	Max Drawdown	Annualized Volatility	Sharpe Ratio
Long(Q1)	14.02%	32.73%	20.64%	0.56
Long(Q1) (with 0.3% Fee)	11.10%	34.22%	20.64%	0.42
CSI 1000 Index	4.84%	46.71%	23.94%	0.10
Q1 Excess Return (vs. CSI 1000 Index)	9.18%	8.28%	16.97%	0.23

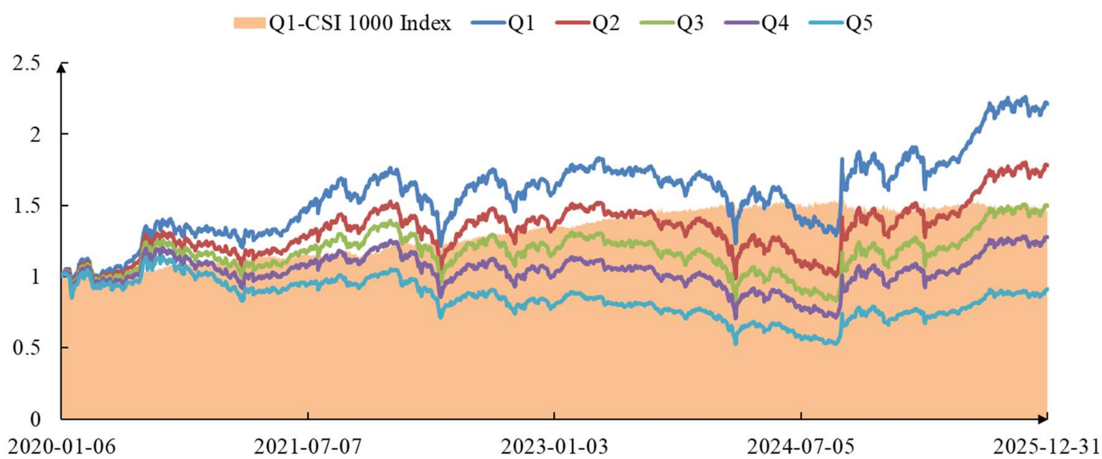

Figure 10: Cumulative Return of the CSI 1000

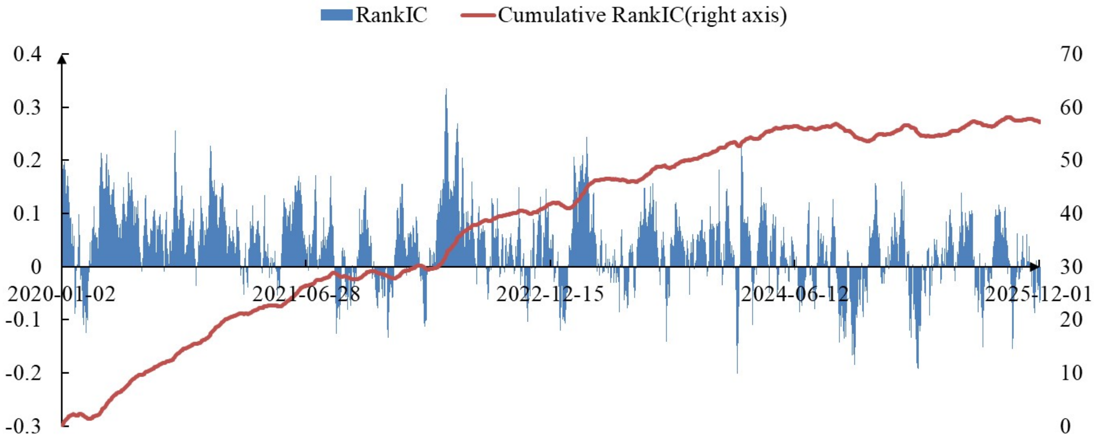

Figure 11: Cumulative Rank IC of the CSI 1000

5.2.5 ChiNext

In the ChiNext sector, Q1 delivers an annualized return of 23.54%, compared with 10.32% for the ChiNext Composite Index, corresponding to an excess return of 13.22%. The Q1 turnover is 74.11%, and the annualized return remains 20.37% after the 0.3% bilateral fee. Figures 12 and 13 summarize cumulative returns and Rank IC; the historical average Rank IC is 6.2%, indicating a stronger rank-order association between predicted upward probability and realized returns in this segment.

Table 6: Performance Statistics of the ChiNext

	Annualized Return	Max Drawdown	Annualized Volatility	Sharpe Ratio
Long(Q1)	23.54%	38.68%	24.99%	0.84
Long(Q1) (with 0.3% Fee)	20.37%	39.05%	24.99%	0.71
ChiNext Composite Index	10.32%	51.08%	27.73%	0.28
Q1 Excess Return (vs. ChiNext Composite Index)	13.22%	23.33%	21.59%	0.28

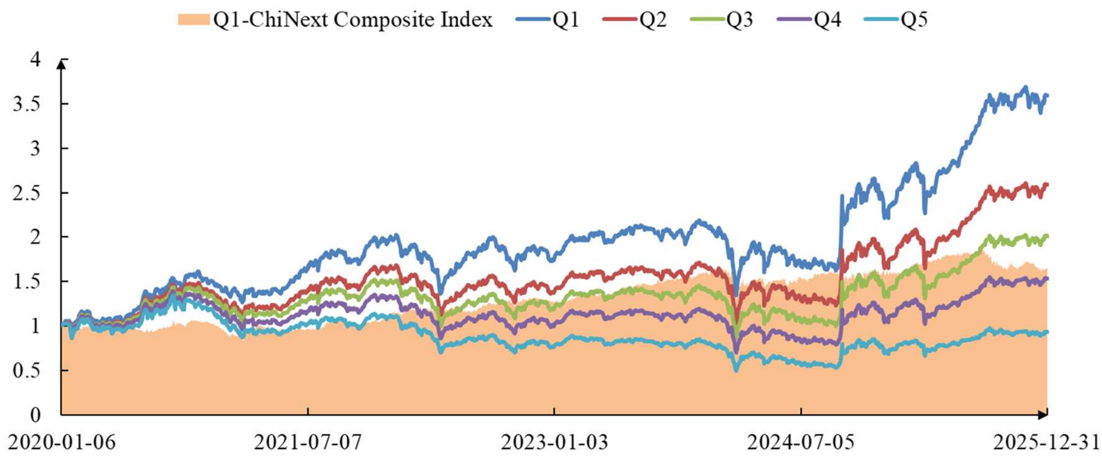

Figure 12: Cumulative Return of the ChiNext

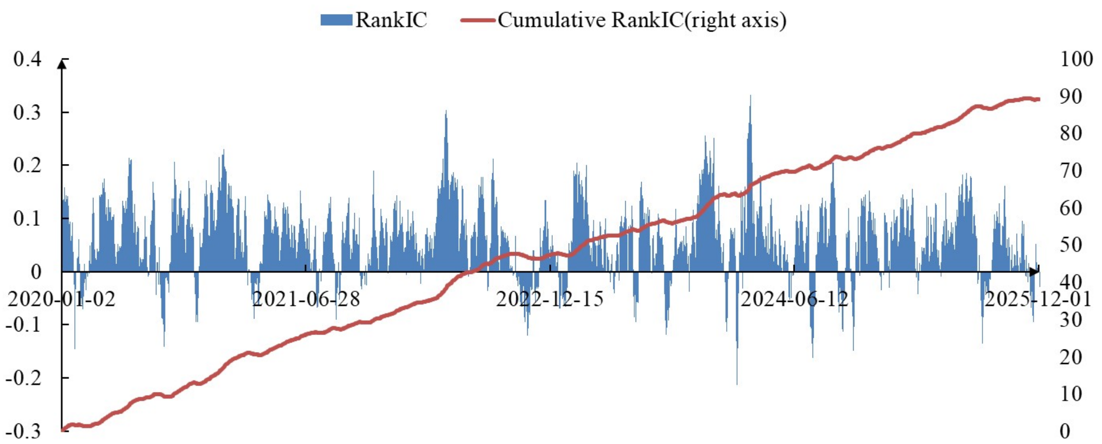

Figure 13: Cumulative Rank IC of the ChiNext

6. Conclusion

Forecasting future stock-price movements from price–volume data has attracted substantial scholarly and practical interest. However, much of the prior literature either models raw price–volume time series directly or relies primarily on K-line chart–based technical analysis, both of which have well-known limitations in capturing richer, morphology-level information embedded in price–volume dynamics.

In this study, we develop a deep-learning framework for stock-price prediction using price–volume information and propose a standardized chart-based methodology. Specifically, starting from raw price–volume sequences, we construct standardized price–volume charts that integrate core market variables (open, high, low, close, and turnover amount) together with widely used derived indicators, including moving averages (MA) and moving average convergence divergence (MACD). We then employ the proposed Efficient Stock Price Prediction Convolutional Neural Network to model the relationship between these standardized charts and subsequent stock-price movements, treating price–volume trend recognition as an image-understanding problem.

We evaluate the model using out-of-sample data from January 2020 to December 2025 across the full Shanghai and Shenzhen A-share market, as well as major sub-sectors (CSI 300, CSI 500, CSI 1000, and ChiNext). On each cross-sectional date, stocks are ranked by the CNN-generated probability of a future price increase, and portfolios are formed by splitting stocks into five equal quintiles in descending order of predicted probability. The results show that the proposed CNN extracts informative signals from standardized price–volume charts, enabling effective identification and classification of stock-price trends. Portfolio returns constructed from these predictions exhibit clear monotonic patterns, and the long portfolios deliver economically meaningful excess returns relative to market benchmarks.

Performance is particularly strong in the full-market and ChiNext settings. In the overall A-share market, the long portfolio achieves an annualized return of 18.20%, compared with 4.34% for the CSI All Share Index, corresponding to an annualized excess return of 13.86%. In the ChiNext sector, the long portfolio achieves an annualized return of 23.54%, versus 10.32% for the ChiNext Composite Index, yielding an annualized excess return of 13.22%.

Overall, this study offers a chart-based deep-learning perspective on stock-price prediction and adds to the financial-technology literature. However, several limitations warrant attention. First, backtesting assumes ideal execution at VWAP; in practice, liquidity shocks and slippage during extreme market conditions could erode realized returns. Second, this study focuses exclusively on the unique microstructure of the China A-share market, and the model's generalizability to other markets remains to be verified.

Future research can address these challenges from two directions: (1) Multimodal Fusion: integrating fundamental data (e.g., financial statements) and unstructured alternative data (e.g., news sentiment) with technical charts to build a more comprehensive predictive framework; (2) Architecture Evolution: exploring advanced vision architectures such as Vision Transformers (ViT) to capture long-range spatiotemporal dependencies that traditional CNNs may overlook.

Acknowledgement: Not applicable

Author contributions: Tao Lin led the overall problem formulation, theoretical framework, and supervision of the study, and contributed to manuscript drafting and revision. Zhuming Chen was primarily responsible for algorithm design, implementation, and experimental evaluation. Ningning An focused on the literature review, comparative analysis, and interpretation of results. Yuanwen Chen handled data collection and preprocessing, supported experimental validation, and assisted in preparing figures, tables, and the final manuscript.

Funding: This work was supported by the National Natural Science Foundation of China (Grant No. 72172164).

Declarations: No Conflict of interest.

See: https://inatgi.in/index.php/jbda/index for more information. DOI:10.63646/j.jbda.2025.030201

References

Abu-Mostafa, Y. S., & Atiya, A. F. (1996). Introduction to financial forecasting. *Applied Intelligence, 6*(3), 205–213. https://doi.org/10.1007/BF00126626

Ballings, M., Van den Poel, D., Hespeels, N., & Gryp, R. (2015). Evaluating multiple classifiers for stock price direction prediction. *Expert Systems with Applications, 42*(20), 7046–7056. https://doi.org/10.1016/j.eswa.2015.05.013

Chen, Y., An, N., & Luo, J. (2023). *Deep learning research report: AI recognition and classification of stock price trends based on convolutional neural networks*. Research Centre and Development of GF Securities.

Dey, S., Kumar, Y., Saha, S., & Basak, S. (2016). *Forecasting to classification: Predicting the direction of stock market price using Xtreme Gradient Boosting*. PESIT South Campus.

Di Persio, L., & Honchar, O. (2017). Recurrent neural networks approach to the financial forecast of Google assets. *International Journal of Mathematics and Computers in Simulation, 11*, 7–13.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance, 25*(2), 383–417. https://doi.org/10.2307/2325486

Guresen, E., Kayakutlu, G., & Daim, T. U. (2011). Using artificial neural network models in stock market index prediction. *Expert Systems with Applications, 38*(8), 10389–10397. https://doi.org/10.1016/j.eswa.2011.02.068

Hu, Y., Liu, K., Zhang, X., Su, L., Ngai, E. W. T., & Liu, M. (2015). Application of evolutionary computation for rule discovery in stock algorithmic trading: A literature review. *Applied Soft Computing, 36*, 534–551. https://doi.org/10.1016/j.asoc.2015.07.008

Imandoust, S. B., & Bolandraftar, M. (2013). Application of k-nearest neighbor (KNN) approach for predicting economic events: Theoretical background. *International Journal of Engineering Research and Applications, 3*(5), 605–610.

Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on Machine Learning* (pp. 448–456). PMLR. http://proceedings.mlr.press/v37/ioffe15.html

Jiang, J., Kelly, B., & Xiu, D. (2023). (Re-)Imag(in)ing price trends. *The Journal of Finance, 78*(6), 3193–3249. https://doi.org/10.1111/jofi.13268

Kingma, D. P., & Ba, J. (2014). *Adam: A method for stochastic optimization*. arXiv. https://arxiv.org/abs/1412.6980

Li, S., Jin, X., Xuan, Y., Zhou, X., Chen, W., Wang, Y. X., & Yan, X. (2019). Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting. *Advances in Neural Information Processing Systems, 32*.

Lim, B., Arık, S. Ö., Loeff, N., & Pfister, T. (2021). Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting, 37*(4), 1748–1764. https://doi.org/10.1016/j.ijforecast.2021.03.012

Nguyen, T. H., Shirai, K., & Velcin, J. (2015). Sentiment analysis on social media for stock movement prediction. *Expert Systems with Applications, 42*(24), 9603–9611. https://doi.org/10.1016/j.eswa.2015.07.052

See: https://inatgi.in/index.php/jbda/index for more information. DOI:10.63646/j.jbda.2025.030201

Pai, P.-F., & Lin, C.-S. (2005). A hybrid ARIMA and support vector machines model in stock price forecasting. *Omega, 33*(6), 497–505. https://doi.org/10.1016/j.omega.2004.07.024

Park, C.-H., & Irwin, S. H. (2007). What do we know about the profitability of technical analysis? *Journal of Economic Surveys, 21*(4), 786–826. https://doi.org/10.1111/j.1467-6419.2007.00519.x

Peng, Y., Liu, Y., & Zhang, R. (2019). Stock price prediction modeling and analysis based on LSTM. *Computer Engineering and Applications, 55*(11), 209–212.

Ren, J., & Wang, A. (2023). Stock price prediction research fusing causal attention Transformer model. *Computer Engineering and Applications, 59*(13), 325–334.

Shah, D., Isah, H., & Zulkernine, F. (2019). Stock market analysis: A review and taxonomy of prediction techniques. *International Journal of Financial Studies, 7*(2), Article 26. https://doi.org/10.3390/ijfs7020026

Tyssedal, J. S., & Tjøstheim, D. (1988). An autoregressive model with suddenly changing parameters and an application to stock market prices. *Journal of the Royal Statistical Society: Series C (Applied Statistics), 37*(3), 353–369. https://doi.org/10.2307/2347310

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems, 30*.

Yang, B., Gong, Z. J., & Yang, W. (2017). Stock market index prediction using deep neural network ensemble. In *2017 36th Chinese Control Conference (CCC)* (pp. 3882–3887). IEEE. https://doi.org/10.23919/ChiCC.2017.8027964

Zhao, H., & Xue, L. (2021). Research on stock prediction based on LSTM-CNN-CBAM model. *Computer Engineering and Applications, 57*(3), 203–207.

Zhong, X., & Enke, D. (2017). Forecasting daily stock market return using dimensionality reduction. *Expert Systems with Applications, 67*, 126–139. https://doi.org/10.1016/j.eswa.2016.09.027

Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., & Zhang, W. (2021). Informer: Beyond efficient transformer for long sequence time-series forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence, 35*(12), 11106–11115. https://doi.org/10.1609/aaai.v35i12.17325

Zhou, Z., & He, X. (2023). Stock price prediction method based on optimized LSTM model. *Statistics and Decision, 39*(6), 143–148.