Wall Street Through a Neural Lens

Forecasting NASDAQ and NYSE Stock Price with Residual Network 2D CNN Model

Sarah Eve Tucker †
School of Computer Science & Digital Technologies
London South Bank University
London United Kingdom
tuckers9@lsbu.ac.uk

ABSTRACT

Stock market forecasting is a complex challenge due to inherent volatility driven by unpredictable economic and geopolitical factors. Traditional models often struggle to capture these dynamics, leading to inconsistent predictions. This study explores the potential of image-based techniques using a residual network-based 2D convolutional neural network (CNN) to improve forecasting accuracy for NASDAQ and NYSE stock prices. By transforming numerical data into grid-like image structures, the model captures spatial, temporal and inter-feature relationships that conventional methods overlook.

Three experiments with varying feature sets demonstrated high accuracy under standard conditions (R^2 up to 0.99), though the model struggled when forecasting future horizons. These findings highlight the promise of image-based forecasting while underlining the need for improved handling of unseen market volatility.

CCS CONCEPTS

• Deep Learning (CNN, Residual Networks) • Financial Market Forecasting • Image-Based Data Transformation • Time-Series Prediction

KEYWORDS

Stock Market Forecasting, Convolutional Neural Network (CNN), Residual Network (ResNet), NASDAQ & NYSE Analysis, Image-Based Stock Prediction

ACM Reference format:

Sarah Eve Tucker. 2025. Wall Street Through a Neural Lens: Forecasting NASDAQ and NYSE Stock Price with Residual Network 2D CNN Model. In *Proceedings of ACM womENcourage 2025, Computer Science: a Catalyst for Educational Change. ACM, Brasov, Romania, 2 pages.*

*Wall Street Through a Neural Lens: Forecasting NASDAQ and NYSE Stock Price with Residual Network 2D CNN Model
†Sarah Eve Tucker

1 Introduction

Stock markets are dynamic systems influenced by a multitude of factors, from economic indicators to investor sentiment. Traditional forecasting methods, such as ARIMA and LSTM, often fail to account for the non-linear and high-dimensional nature of financial data. This study investigates whether transforming numerical stock data into image-like representations improves predictive accuracy for NASDAQ and NYSE prices using a residual network 2D CNN.

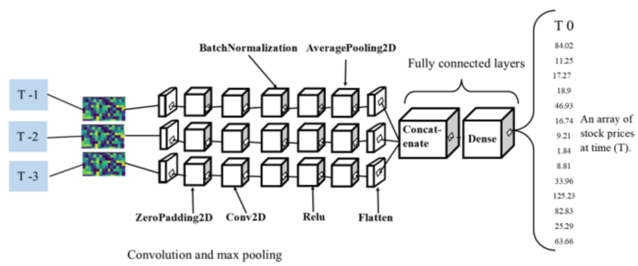

Figure 1: Residual network architecture: Temporal stock data (T_{-1} to T_{-3}) processed through convolutional and pooling layers, with skip connections, to predict prices at T_0.

2 Methodology

2.1 Data Transformation

Numerical stock data was transformed into images using a technique termed Grid-Based Temporal-Feature Structuring (GTFS) technique to predict NASDAQ and NYSE stock prices. GTFS encodes weekly stock prices and economic indicators into 2D grids, enabling the CNN to capture spatial and temporal patterns [1] & [2]. The dataset includes 3,280 weekly stock prices from June 2020 to March 2024, alongside macroeconomic indicators such as CPI, industrial production and money supply. Data preprocessing included Min-Max scaling to normalise features [3] and Pearson's Correlation to select relevant variables [4]. The model architecture combines convolutional layers, residual connections and max pooling to extract hierarchical patterns, with dropout layers (10%) to mitigate overfitting [5]. The Adam optimiser was used for training, with learning rates adjusted to improve convergence.

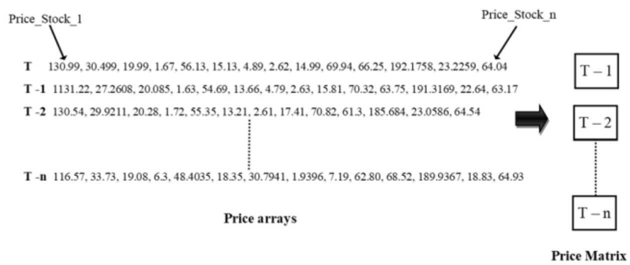

Figure 2: Price array transformation into structured matrix: Sequential stock prices (T, T_{-1}, T_{-2}, ... T_{-n}) organised into rows for model input.

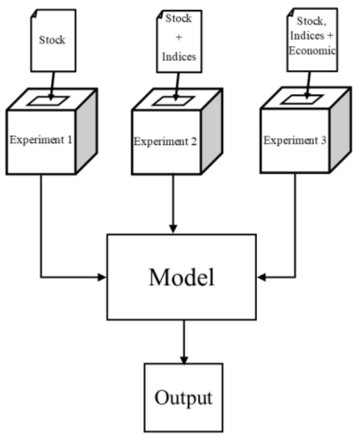

Figure 3: The integration of three experiments into a single model leading to an output.

2.2 Evaluation Metrics

The data was split into 80% training and 20% validation sets. Models were trained over 175 epochs and evaluated using Mean Squared Error (MSE), R^2 score and volatility analysis. Generalisability was assessed through future horizon forecasting.

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)2}{\sum_{i=1}^{n}(y_i - \bar{y}_i)2} \quad (1)$$

R^2 Score formula for evaluating prediction accuracy.

$$MSE = \frac{1}{n}\sum_{i=1}^{n}\left(y_i - \hat{y}_i\right)2 \quad (2)$$

MSE formula for measuring average prediction error.

3 Results

The experiments demonstrated progressively stronger predictive accuracy, with E1 (historical prices only) achieving an R^2 of 0.97, E2 (enhanced with correlated indices) reaching 0.981, and E3 (incorporating macroeconomic indicators) attaining a near-perfect 0.99. However, the marginal improvement (0.009 over E2) suggests diminishing returns from integrating supplementary features.

While standard test conditions produced high R^2 scores across all experiments, persistent outliers highlight ongoing volatility challenges. Horizon forecasting results for both E2 and E3 revealed the model's difficulty in effectively utilising the extra information to aid prediction accuracy on unseen future data. This performance dip when extending beyond trained periods indicates that, although additional features improve accuracy under stable conditions, their contribution is less meaningful in volatile or unfamiliar contexts.

	E1	E2	E3
R^2 Score	0.97	0.981	0.99

Table 1: R^2 scores for standard testing across three experiments.

4 Discussion

The residual network's ability to capture spatial-temporal patterns in image-transformed data outperformed traditional numerical methods. The high R^2 scores in standard test conditions confirm the model's strength in identifying relationships within the images. However, its performance dips during extreme volatility for E2 and E3, indicating a need for adaptive mechanisms. This suggests that while the model benefits from historical price trends, it struggles to extract useful information from supplementary features in unseen scenarios. Future research could explore incorporating daily data to capture intra-week fluctuations to address volatility gaps identified in horizon forecasting, alongside refining feature engineering to reduce noise and improve generalisability.

ACKNOWLEDGMENTS

The author thanks Dr. Ali Salimian and Francis Babayemi for their invaluable guidance and the School of Computer Science and Digital Technologies for providing essential resources.

REFERENCES

[1] G. Anand and R. Nayak, "DeLTa: Deep local pattern representation for time-series clustering and classification," Knowledge-Based Systems, vol. 212, 2020.

[2] S. Barra, S. M. Carta, A. Corriga, A. S. Podda and D. R. Recupero, "Deep Learning and Time Series-to-Image Encoding for Financial Forecasting," CAA Journal of Automatica Sinica, vol. 7, no. 3, pp. 683-692, 2020.

[3] S. Jeong, I. Park, H. S. Kim, C. H. Song and H. K. Kim, "Temperature Prediction Based on Bidirectional Long Short-Term Memory and Convolutional Neural Network Combining Observed and Numerical Forecast Data," Sensors, vol. 21, no. 3, pp. 1-20, 2021.

[4] G. Li, A. Zhang, Q. Zhang, D. Wu and C. Zhan, "Pearson Correlation Coefficient-Based Performance Enhancement of Broad Learning System for Stock Price Prediction," IEEE, vol. 69, no. 5, pp. 2413 - 2417, 2022.

[5] M. Kirisci and O. C. Yolcu, "A CNN-based model for financial time series forecasting: TAIEX and FTSE Stocks Forecasting," Neural Processing Letters, vol. 54, no. 4, pp. 3357 - 3374, 2022.