ADDRESSING FINANCIAL MARKET UNCERTAINTY: EXTREMA PREDICTION WITH SEQUENCE MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Deep learning models, particularly sequence-based architectures, are widely used for trend prediction and time series analysis in financial markets. This paper investigates a fundamental aspect of chart pattern formations: the prediction of local extrema types. By classifying extrema into four distinct categories using historical extrema data, we aim to provide a novel perspective on chart pattern identification. However, our findings reveal that these models struggle to generalize under data distribution shifts, achieving significantly lower prediction accuracy on out-of-training data. These results underscore the limitations of deep learning-based strategies in dynamic financial environments and highlight the need for robust methods to address market variability.

1 INTRODUCTION

Patterns in OHLC (Open, High, Low, Close) candlestick charts are recurring sequences that traders have historically used as signals for buy-and-sell decisions. In technical analysis, these patterns, including head-and-shoulders, double tops, and triangles, are often combined with other tools to inform trading strategies. Many of these patterns can be identified through the analysis of extrema—local maxima and minima in price data.

Extrema are broadly categorized into higher highs, lower lows, higher lows, and lower highs, based on their position relative to preceding extrema. Highs represent local maxima, while lows are local minima. These extrema provide critical insights into price trends, with higher highs and lower lows defining upward or downward movements. Accurate prediction of extrema types allows traders to anticipate price rises and falls, facilitating earlier market entry and potentially maximizing gains. Deep learning models often achieve low mean squared error (MSE) in predicting future prices, but this can mask their inability to accurately predict trends Wang et al. (2018), Kaya et al. (2020). Although the loss may be low, these models frequently fail to identify whether prices will rise or fall,and often exploit superficial correlations in historical price data rather than capturing underlying market trends.

While previous studies have explored machine learning models to identify extrema, emphasizing their significance, they have largely overlooked the prediction of future extrema types Sokolovsky & Arnaboldi (2020). This paper investigates whether deep learning models, particularly sequence-based architectures, can accurately classify the type of the next two extrema based on historical extrema data. These predictions are tested against future data to assess the model's ability to generalize.

2 METHODOLOGY

We experimented with three distinct sequence models— Atthention based GRU, hybrid CNN+LSTM, and ARIMA—to predict and classify the types of the next two price extrema. The S&P 500 and NIFTY 50 indices were chosen as representative datasets due to their diverse market behaviors. The S&P 500 includes 500 large-cap companies from various sectors in the United States, while the NIFTY 50 represents the top 50 companies across industries on the National Stock Exchange of India. Daily OHLC (Open, High, Low, Close) prices were collected for a two-year period, providing sufficient data for robust training and evaluation.

Extrema were identified using a sliding window of size 5, ensuring that each extrema had at least five neighboring points that were lower (for highs) or higher (for lows). This approach filtered out noise, capturing only meaningful price movements. Each identified extrema was labeled into one of four categories: higher high (HH), lower low (LL), higher low (HL), or lower high (LH), based on its relative position compared to previous extrema. After identifying the extrema, the training data is constructed by creating sequences of 5 consecutive extrema, which serve as input features along with trend indicator values at these points, and predicting the next 2 extrema as the target output (y-values)

To improve the predictive performance of our models, we integrated additional features commonly used in technical analysis, including MACD, RSI, and Bollinger Bands. These indicators provided supplementary insights into market trends and volatility, enhancing the dataset's richness. The attention-based GRU model combines the power of both attention mechanism and sequence modelling(Jung et al. (2021)). The hybrid CNN+LSTM model combined convolutional layers to detect local patterns with LSTM layers to model long-term dependencies(Liu et al. (2017)). Finally, the ARIMA model served as a baseline, focusing on linear time-series forecasting(Ho & Xie (1998)). Each model was trained to predict the types of the next two extrema based on historical extrema data and the additional technical indicators. The trained models were then evaluated on unseen future data corresponding to each index.

3 RESULTS

As observed in all different sequence models across different stocks, while the training accuracy reaches a high of 98%, the accuracy of predicting the next two extrema correctly averages around 75%, which is insufficient for making reliable trading decisions (Table 1). The sequence models were not accurately able to analyze the alternative and recursive sequence of highs and lows. Moreover, even when considering the probability of predicting the type of at least one extrema accurately, the performance remains not upto the mark.

Table 1: Train and Test Accuracies for All models(%)

Model	**Index**	**Train Acc.**	**Test Acc. (Both)**	**Test Acc. (Atleast 1)**
ARIMA	S&P 500	98.73	74.32	84.86
	NIFTY 50	98.05	76.54	83.61
GRU	S&P 500	97.45	73.12	82.78
Attention	NIFTY 50	97.89	75.33	82.12
CNN+	S&P 500	96.89	72.57	80.91
LSTM	NIFTY 50	97.12	74.89	81.03

4 CONCLUSION

This paper highlights the limitations of sequence-based deep learning models in predicting future price extrema, such as maxima and minima, and their struggle to accurately classify extrema types. These models fail to capture the underlying dynamics of financial markets, especially when confronted with real-world data distribution shifts. The results emphasize the need to integrate additional factors—such as sentiment, behavioral biases, and market psychology—into predictive models. Furthermore, the paper critiques the reliance on conventional performance metrics, like mean squared error, in financial applications. While models may show low error rates in price predictions, they fall short in capturing market trends and extremas, which are vital for making informed trading decisions. In conclusion, deep learning models require a more comprehensive approach in financial forecasting—one that incorporates both price data and behavioral insights to improve accuracy and robustness.

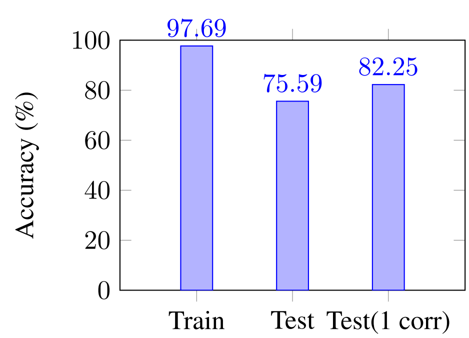

Figure 1: Comparison of Average Training and Test Accuracies for NIFTY 50

REFERENCES

S.L. Ho and M. Xie. The use of arima models for reliability forecasting and analysis. *Computers Industrial Engineering*, 35(1):213–216, 1998. ISSN 0360-8352. doi: https://doi.org/10.1016/S0360-8352(98)00066-7. URL https://www.sciencedirect.com/science/article/pii/S0360835298000667.

Seungmin Jung, Jihoon Moon, Sungwoo Park, and Eenjun Hwang. An attention-based multilayer gru model for multistep-ahead short-term load forecasting. *Sensors*, 21(5), 2021. ISSN 1424-8220. doi: 10.3390/s21051639. URL https://www.mdpi.com/1424-8220/21/5/1639.

Celal Kaya, Alperen Yılmaz, Gizem Uzun, and Zeynep Kilimci. Stock pattern classification from charts using deep learning algorithms. *Academic Perspective Procedia*, 3:445–454, 10 2020. doi: 10.33793/acperpro.03.01.89.

Shuanglong Liu, Chao Zhang, and Jinwen Ma. Cnn-lstm neural network model for quantitative strategy analysis in stock markets. pp. 198–206, 10 2017. ISBN 978-3-319-70095-3. doi: 10.1007/978-3-319-70096-0_21.

Artur Sokolovsky and Luca Arnaboldi. Machine learning classification of price extrema based on market microstructure features: A case study of sp500 e-mini futures, 09 2020.

Jia Wang, Tong Sun, Benyuan Liu, Yu Cao, and Degang Wang. Financial markets prediction with deep learning. In *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pp. 97–104, 2018. doi: 10.1109/ICMLA.2018.00022.

A APPENDIX

DATA PREPROCESSING

Training data was obtained for the period between 2016 and 2018 to avoid the market volatility and unpredictability associated with the COVID-19 pandemic. The model was then tested on data from the 2019 period to evaluate its performance in a more stable market environment. We identify extrema points in the normalized closing price data. For maxima, a sample is considered an extrema if its neighbors on both sides have lower amplitudes. Similarly, for minima, the sample must have a lower amplitude than its neighbors. This logic is implemented efficiently using a linear search algorithm. In cases where flat peaks or valleys occur—where several consecutive samples have the same amplitude—the middle sample is selected as the extrema for consistency. After identifying the extrema, the training data is constructed by creating sequences of 5 consecutive extrema, which serve as input features along with trend indicator values at these points, and predicting the next 2 extrema as the target output (*y-values*). The extrema are labeled into one of the following four categories:

1. **HH (Higher High)**: This category is assigned when the current extremum is a higher high, meaning it is a local maximum where the price is greater than the previous maxima.
2. **LL (Lower Low)**: This category is assigned when the current extremum is a lower low, meaning it is a local minimum where the price is lower than the previous minima.
3. **LH (Lower High)**: This category is assigned when the current extremum is a lower high, meaning it is a local maximum where the price is lower than the previous high.
4. **HL (Higher Low)**: This category is assigned when the current extremum is a higher low, meaning it is a local minimum where the price is higher than the previous low.

FEATURE ENGINEERING

As stated, we enhanced the dataset with additional engineered features derived from the OHLC (Open, High, Low, Close) data , obtained from yahoo finance [1].. These features include:

[1] https://finance.yahoo.com

- **MACD (Moving Average Convergence Divergence)**: Measures the difference between the short-term (typically 12-period) and long-term (typically 26-period) exponential moving averages (EMAs), providing insight into momentum and trend direction.
- **RSI (Relative Strength Index)**: A momentum oscillator that compares the magnitude of recent price gains to losses over a specified period, offering a measure of overbought or oversold conditions.
- **Bollinger Bands**: Calculated using a moving average and standard deviations, these bands provide information about price volatility and potential breakout points.

Chart Patterns using Extrema Formation

We analyzed the formation and functionality of various chart patterns across different market types, focusing on commonly used patterns such as Double Top and Bottom, Head and Shoulders, Inverse Head and Shoulders, Ascending Triangle, and Descending Triangle. These patterns are essentially permutations of higher, lower, or same-level extrema over a series of points. By defining these patterns in terms of extrema relationships, we implemented their identification in Python. Below is a summary of how these patterns were structured:

1. **Double Top and Bottom**: A Double Top is formed by four extrema, where:

$$B > A, \quad B > C, \quad C > A, \quad D > C$$

   Additionally, extrema B and D are on the same level.

2. **Ascending and Descending Triangles**: These patterns require four extrema points along with a final breakout point where trades are executed. For Ascending Triangles:

$$A > B, \quad B < C, \quad C > D, \quad D > B$$

   Here, A and C are at the same level. This pattern represents a higher high, followed by a higher low, then a same-level high, and another higher low.

3. **Head and Shoulders / Inverse Head and Shoulders**: These patterns are confirmed with five extrema. For a Head and Shoulders pattern:

$$C > A, \quad C > E, \quad A > B, \quad E > D$$

   Additionally, extrema B and D lie at the same level.

The above highlights the sufficiency of correctly identifying future extrema for identifying these chart patterns before hand and maximizing gains.