

Technical Indicator-Based Intelligent Financial Market Forecasting Using Deep Learning

Prem Kumar Gupta

M. Tech, Computer Science and Engineering
Department of Computer Science and Engineering Sarvepalli Radhakrishnan University, Bhopal,
India

Dr. Varsha Namdeo

Professor, Department of Computer Science and Engineering Sarvepalli Radhakrishnan University,
Bhopal, India

Abstract

1. Introduction

Financial markets pump out a staggering amount of data every day, both old and new. Investors want any edge they can get to recognize trends and dodge losses—but let's be real, the usual forecasting tools can't keep up with all the market's quirks and wild swings.

Technical indicators make things simpler. They take all those past price changes and turn them into signals you can actually use—things like spotting a trend, catching shifts in momentum, flagging volatility spikes, or sending up a warning when a reversal might be coming. Pair them with deep learning models, and now you've got a forecasting system that's a lot sharper and doesn't break down so easily.

This research digs into how building features from technical indicators can actually make deep learning models better at predicting what happens in the market.

PROPOSED METHODOLOGY

3.1 Introduction

Every trading day, financial markets spit out a huge amount of data, minute by minute. Prices jump, trading volume shifts, volatility spikes, and investor sentiment shifts. But raw market data isn't perfect—it's noisy, full of missing values, and packed with messy, nonlinear patterns that make prediction tough. If you want any hope of forecasting market moves, you need a solid plan for turning all that raw information into useful signals first.

This approach brings together technical indicators and deep learning to boost forecasting accuracy. Here’s how it works: First, you pull in historical market data. Then, you clean and prep it. After that, you generate technical indicators and build features that matter. Next, you train a deep learning model. Finally, you judge how well the model predicts using common statistical measures.

The process breaks down into six main steps:

- 1. Collect the data
- 2. Clean and prepare it
- 3. Generate technical indicators
- 4. Engineer features for your model
- 5. Train your deep learning model
- 6. Evaluate how well it predicts

With each phase, you cut down on errors and sharpen your predictions.

3.2 Overall Research Framework

The proposed framework is illustrated in Figure 3.1.




Figure 3.1: Proposed methodology for financial market forecasting.

3.3 Data Collection

Everything starts with gathering old financial market data. These records are the backbone for figuring out how the market actually behaves and spotting trends you might use to predict where prices are headed next. Here, the dataset pulls together daily market stats from places like Yahoo Finance—each entry breaks down various details from a single day’s trading session.

Table 3.1 Historical Dataset Description

Attribute	Description
Date	Trading Date
Open	Opening Price
High	Highest Price
Low	Lowest Price
Close	Closing Price
Volume	Trading Volume

Historical data covering several years is collected to ensure sufficient variability in market conditions, including bullish, bearish, and sideways trends.

3.4 Data Preprocessing

Raw financial data is messy. Gaps, weird values, and repeats can throw off any prediction model, so cleanup comes first.

Here’s how the preprocessing works:

Missing values get flagged right away. Sometimes, you just toss incomplete records. Other times, you fill the gaps with averages or whatever the data calls for.

Duplicates are a pain, too. If the same record shows up more than once, you delete the extras. That way, the model doesn’t learn the same thing twice.

Outliers stand out, and not in a good way. You spot extreme values—maybe someone made a typo or the instruments glitched. You need to check them before feeding data to your model.

Normalization is last. Financial data can have numbers all over the place, which messes with models. Min-Max normalization pulls every value onto the same scale so each feature carries equal weight.

And here’s the Min-Max normalization equation:


$$X_{\text{norm}} = \frac{X - X_{\min}}{X_{\max} - X_{\min}}$$

where:

- XXX = Original Value
- $X_{minX_{\{min\}}}$ = Minimum Value
- $X_{maxX_{\{max\}}}$ = Maximum Value

Normalization accelerates model convergence and improves prediction stability.

Figure 3.2 Data Preprocessing Pipeline



3.5 Technical Indicator Generation

Technical indicators take past price data and turn it into signals that tell us about market momentum, trends, and volatility. This framework pulls out several of these indicators from historical market data.

Table 3.2 Technical Indicators

Indicator	Purpose	Interpretation
RSI	Momentum	Overbought/Oversold
MACD	Trend	Buy/Sell Signals
SMA	Trend	Long-Term Average
EMA	Trend	Recent Price Trend
ATR	Volatility	Market Risk
Bollinger Bands	Volatility	Price Expansion & Contraction

Relative Strength Index (RSI)

RSI tracks how fast and how much prices have been changing lately. Most traders read it like this:

- RSI above 70: The market's probably overbought.
- RSI below 30: The market's considered oversold.

Moving Average Convergence Divergence (MACD)

MACD uses exponential moving averages to spot which way the market's trending. It throws up buy or sell signals when the moving average lines cross.


Bollinger Bands

Bollinger Bands consist of:

- Upper Band
- Middle Moving Average
- Lower Band

These bands expand during high volatility and contract during low volatility.

Figure 3.3 Technical Indicator Framework



3.6 Feature Engineering

Feature engineering combines raw price variables with generated technical indicators to create a richer representation of market behavior. Instead of relying only on historical closing prices, the proposed framework uses a multidimensional feature set including:

- Open Price
- High Price
- Low Price
- Close Price
- Volume
- RSI
- MACD
- SMA
- EMA
- ATR
- Bollinger Bands

This expanded feature space allows the deep learning model to learn more complex relationships within the financial time series.

3.7 Chapter Summary

In this chapter, I laid out the approach for forecasting financial markets using technical indicators. First, I start by gathering and cleaning the data. Once that's done, I generate technical indicators and engineer new features. These improved features go straight into a deep learning prediction model. To check how well the method works, I use standard performance metrics. This step-by-step process offers a solid way to boost both the accuracy and reliability of market forecasts. Thanks for all the support during this research.

REFERENCES

1. S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
2. K. Cho et al., "Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation," *EMNLP*, pp. 1724–1734, 2014.
3. I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.
4. D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," *ICLR*, 2015.
5. A. Vaswani et al., "Attention Is All You Need," *NeurIPS*, 2017.
6. B. Lim and S. Zohren, "Time-Series Forecasting with Deep Learning: A Survey," *Philosophical Transactions of the Royal Society A*, vol. 379, 2021.
7. J. Murphy, *Technical Analysis of the Financial Markets*. New York Institute of Finance, 1999.
8. S. Achelis, *Technical Analysis from A to Z*. McGraw-Hill, 2001.
9. Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," *Nature*, vol. 521, pp. 436–444, 2015.
10. A. Géron, *Hands-On Machine Learning with Scikit-Learn, Keras & TensorFlow*, 3rd ed. O'Reilly Media, 2022.
11. G. James, D. Witten, T. Hastie, and R. Tibshirani, *An Introduction to Statistical Learning*, Springer, 2021.
12. F. Chollet, *Deep Learning with Python*, 2nd ed. Manning Publications, 2021.
13. R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. MIT Press, 2018.
14. J. Han, M. Kamber, and J. Pei, *Data Mining: Concepts and Techniques*, 3rd ed. Morgan Kaufmann, 2011.
15. T. Young, D. Hazarika, S. Poria, and E. Cambria, "Recent Trends in Deep Learning Based Natural Language Processing," *IEEE Computational Intelligence Magazine*, vol. 13, no. 3, pp. 55–75, 2018.
16. M. Schuster and K. K. Paliwal, "Bidirectional Recurrent Neural Networks," *IEEE Transactions on Signal Processing*, vol. 45, no. 11, pp. 2673–2681, 1997.
17. J. Brownlee, *Deep Learning for Time Series Forecasting*. Machine Learning Mastery, 2018.
18. Y. Li and X. Zhao, "Comparative Analysis of LSTM and GRU Models for Stock Market Prediction," *IEEE Access*, 2024.
19. J. Wang et al., "Financial Forecasting with Hybrid Deep Learning Architectures," *Expert Systems with Applications*, 2024.
20. M. Qiu, Y. Song, and F. Liu, "Stock Price Prediction Using Deep Learning Models," *IEEE Access*, 2020.