

AN INTERPRETABLE DEEP LEARNING BASED DECISION SUPPORT SYSTEM FOR STOCK FORECASTING AND INVESTMENT RECOMMENDATION USING TEMPORAL FUSION TRANSFORMER

D. MANJU ¹, M. KOTESWARA RAO ², RAJESH KUMAR VERMA ³, PADMINI DEBBARMA ⁴, ANAND KUMAR SARASWATHI RATHOD ⁵

¹ Assistant Professor, Department of CSE-(CyS,DS) and AI&DS, Hyderabad, India.

² Assistant Professor, Department of Information Technology, VNR Vignana Jyothi Institute of Engineering and Technology, Hyderabad, Telangana, India.

³ Professor, Department of CSE, CMR College of Engineering, Hyderabad, India,

⁴ Ph.D Scholar, ICFAI, Hyderabad, India, debbarmapadmini.scholar24@ifheindia.org

⁵ Professor of Practice, Department of CSE-(CyS,DS) and AI&DS, Vallurupalli Nageswara Rao Vignana Jyothi Institute of Engineering and Technology, Hyderabad, India, sr_ak@yahoo.com

E-Mail : nuthana525@gmail.com , mailtomkraj@gmail.com, rajeshverma.hydr10@gmail.com , debbarmapadmini.scholar24@ifheindia.org, sr_ak@yahoo.com

ABSTRACT

The stock market has rapid changes caused by economic conditions, company performance and investor behavior. Because of this, stock price predictions and correct investment decisions are hardly achieved. Investors today use computer-based systems to study large volumes of financial data and find useful patterns barely perceived using traditional methods. However, most existing stock prediction and recommendation systems that have some major problems. Many systems use only technical indicators or only fundamental data, while not both together. Some models cannot properly capture both the short-term and long-term price movements of the stock. Moreover, many of these deep learning models behave like black boxes and fail to explain clearly why any stock is recommended to be on buy, sell or hold. This therefore compromises trust in their prediction capabilities, especially those specifically for highly volatile markets. In this paper, we try to solve those problems by proposing a new stock forecasting and recommendation system based on the Temporal Fusion Transformer (TFT), combining historical stock prices, technical indicators and fundamental financial data in one framework. It can learn both short-term and long-term trends and show which factor is most important for the prediction. After predicting the price and other financial indicators, our model extends TFT into a decision-support system by embedding financial ratio analysis and recommendation logic within the forecasting pipeline.

Keywords: *Deep Learning, Investment Recommendation, Stock Forecasting Temporal Fusion Transformer, Technical Indicators, Fundamental Analysis.*

1 INTRODUCTION

The stock market is significant in the contemporary economy. It is used by companies to raise funds and investors attempt to increase their wealth. The prices of stocks of publicly traded companies are continuously bought and sold on different exchanges every single day. Such price alterations are based on economic patterns, the profitability of the company, government interventions and even investor attitude at that time.

The stock price is however not simple to predict. Prices may either increase or decrease rapidly due to factors such as inflation, increasing interest rates, political instability or world events. More confusion is added by fast trading by computers, market rumors and speculation. Information is not always disseminated equally among investors hence giving some players an edge. Such emotions as the fear of loss, overconfidence, or just being influenced by the crowd may result in poor choices. Investors attempt to cope with such uncertainty by employing such tools as fundamental analysis, which examines the financial status of the company and technical analysis, which examines stock price patterns. Although these

traditional methods may be useful, they may not adequately reflect the complexity of the market in today days.

1.1 Fundamental Analysis

Fundamental analysis is one of the primary ways of determining the actual value of a stock. This is done by examining the financial health and performance of a company by looking at key financial statements such as balance sheet, income statement and cash flow statement. These reports provide significant information regarding such performance parameters as revenue growth, profitability, earnings per share (EPS), total assets and liabilities. With such factors, investors will be in a better position to analyze the long term stability and growth of a company.

Fundamental analysis is particularly applicable to long-term investors who are interested in underpriced companies that have good finances. Nevertheless, it tends to overlook short-term fluctuations in the market, because of news or sentiment, or because of the momentum. This renders it less applicable to short-term traders and day traders who are more dependent on instant price changes and current market trends rather than on the analytical financial data.

1.2 Technical Analysis

Technical analysis, unlike fundamental analysis, is used to forecast future stock prices using previous price changes and volume. It applies the use of indicators such as SMA, EMA, Bollinger Bands, RSI, MACD, stochastic oscillators to demonstrate trends, momentum and possible price reversals. The chart patterns and candlestick patterns are also analyzed to locate the entry and exit points. The technique is used in short term and trend trading. Nevertheless, it can create misleading signals when the market is volatile and frequently fails to consider the economic or company specific aspects. Hence, it may not be as reliable when applied on its own.

1.3 Machine Learning and Deep Learning in Stock Forecasting

Over the last few years, machine learning (ML) and deep learning (DL) have emerged as potent instruments to forecast the prices of stock and give investment advice. ML and DL enable computers to get trained based on past information without specification. Rather than using the human judgment alone, these models are capable of

working through vast quantities of data such as stock prices, news, financial statements and technical analysis so as to determine the way a price might move in the future. The use of machine learning algorithms (decision trees, support vector machines, and random forests) can reveal some useful trends in previous data and aid in deciding whether a stock will increase or decrease. A more advanced type of ML is known as deep learning, which uses models, such as neural networks, which work in a similar way as the human brain. They are time sensitive models and particularly models used to predict stock trends, detect hidden trends and enhance accuracy of recommendations are time-series models which are capable of identifying complex patterns over time.

1.4 Temporal Fusion Transformer (TFT) Model

The Temporal Fusion Transformer (TFT) is a powerful deep learning model built for analyzing time

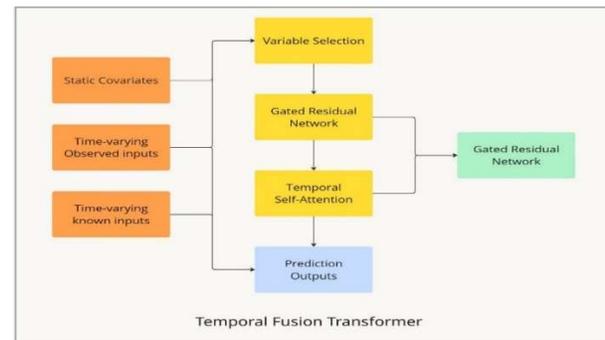


Figure 1: Temporal Fusion Transformer model workflow

The architecture of a Temporal Fusion Transformer (TFT) is depicted in figure 1: the model is an interpretable neural network that is based on multi-horizon time-series forecasting. The model is combined with three forms of inputs, i.e. static covariates, time-varying observed inputs and time-varying known inputs. These inputs are then fed to a Variable Selection module that selects the most pertinent features to the task of forecasting. The chosen variables are then sequenced with a Gated Residual Network (GRN) which performs non-linear transformation whilst maintaining information flow to avoid overfitting and avoid loss of interpretability.

The transformed signals are also optimized by using a Temporal Self-Attention to allow the model to learn various significant time dependencies and to learn long-range temporal correlation. The learned representations are then refined by a secondary GRN which then produces the final multi-step Prediction Outputs. Altogether, TFT architecture is a

combination of competitive selection of variables, gating mechanisms and attention layers that offer precise forecasts maintaining transparency and interpretability.

The primary interest of this study lies in addressing three significant issues that arise in currently available stock forecasting and investment recommendation systems, which work inaccurately regarding the following three aspects: the prediction horizon, ambiguous model decision-making, and usefulness of predictions for real-time stock markets. Currently available stock forecasting systems based on deep learning techniques, including LSTM, CNN-LSTM, GRU, etc. are mostly focused on providing predictions for only the short-term stock prices. Their ability to demonstrate an understanding of long-term stocks trends or sharp changes in stock markets, which play a significantly important role in stock investment, is negligible. Thus, there is a need to see whether Temporal Fusion Transformers, which are proposed for multi-time horizon forecasting, provide more accurate predictions than deep learning-based forecasting techniques for stock prices at short-term, medium-term, or long-term horizons. Therefore, this research will seek to answer the following research questions:

Can the Temporal Fusion Transformer (TFT) model improve the forecast of stock prices of various horizons compared to the deep models used as well as Recurrent Deep Learning Models?

How could the principles of interpretability be integrated into the deep learning networks to predict stocks so that investors can see what predictions were obtained?

Will the proposed framework which involves the TFT-based hybrid result in credible investment decisions of different stocks when the markets are operating in real time conditions?

2 RELATED WORK

The use of Deep Learning Models and hybrid approaches is combating the inaccuracy and conceptual inadequacy of financial forecasting in the stock market with the current dynamic world of financial forecasting. The time-series forecasting tasks have been widely done using the traditional forecasting methods, including the ARIMA model. What is apparent about these methods, however, is that they cannot identify the non-linear character of financial data therefore,

make them less effective in understanding financial activities within the stock market. The prediction of stock prices using the ARIMA model used by [1] established the fact that the traditional time-series forecasting models can also follow short-term financial transactions in the financial market. The model though does not work because it does not consider non-linearities and linearity at the expense of financial data which has high volatility. The idea of the Temporal Fusion

Transformer (TFT) was presented in [2], as a conceptual Deep Learning Model of multi-horizon time-series forecasting tasks to allow higher accuracy and explainability. It was able to use attention mechanisms, feature selection networks, and gating networks all together but failed to examine how it was used to provide recommendations within a financial stock market system. The appropriateness of TFT to the stock price prediction problem was particularly used in [3] and it was demonstrated that the transformer-based architecture is more efficient than the conventional LSTM-based and CNN-based models. This research also largely concentrated on the prediction of the stock prices and did not include the subject of stock investment suggestion along with the integration of the underlying information regarding finances. To enhance the idea of forecasting, [4] used CEEMD-based signal decomposition to eliminate the noise that exists in the financial data. This research enhanced the forecasting performance, but limited it to predictions of price only and prevented addressing the areas of decision-making and basic finances. The recommendation system proposed by [5] was created with the aim of investment planning according to machine learning methods. This paper gave hints of both buying and selling but limited itself to the use of the technical indicators and did not delve into the aspects of basic finances and the decision-making concept. The application of the concept of deep learning methods including LSTM and CNN were implemented by [6] to use them in the process of stock forecasting and pairs trading strategies. The issue of the superiority of the use of the neural network method over the traditional methods was discovered in the obtained results. The system of recommendations including technical trading rules was proposed by [7]. Though this employed a useful method with intraday trading, this was not able to accommodate changes in market trends and dependencies. A machine learning stock investment-recommendation platform was introduced by [8]. It was not the best tool to use as a decision-support tool but it was limited by traditional machine learning and could not do deep learning. The study of [9] examined

the use of classifier-based recommendation systems in Indian stock markets and found that the ensemble models improved the prediction. Nevertheless, they were not combined financial models, and interpreted similarly to problems. A systematic review of technical and fundamental analyses was described by [10] and found that source-specific models were inferior to combined ones. Nevertheless, the majority of reviewed mixed analysis was superficial models. Both technical and fundamental stock market factors were integrated with a multi-layer perceptron model by [11] through the combination of the two. Their findings indicated that their model increased the accuracy of prediction. Nevertheless, the multi-layer perceptron model exhibit issues when it comes to following the dependence of the sequences. [12] examined both the technical and fundamental market factors and came to the conclusion that both the market factors have the insights of complementarity. They however did not come up with a deep learning architecture that exploits these complementarity pieces of information. A system of stock trading and recommendation was proposed by [13] based on the deep learning models. Even though their system boosted trading and investment decision-making, it is not transparent. A stock recommendation system, [14], was developed using cognition. It was a copycat of investor action but failed to apply deep learning in the long term. [15] introduced a stock recommender system, which is an association rule mining of portfolios. It is also not applicable in price prediction in dynamically changing markets, although it is very effective in pattern analysis. The system of stock recommendation offered by [16] was developed on the basis of Elliott Wave Theory and multi-agent technology. It is founded on existing regulations but is weak on adapting in complex markets. A set of stocks prediction and recommendation was offered by [17].

The ensemble learning approach boosted predictiveness and made the model more difficult to interpret and more complex. An example of hybrids was given in [18] in terms of stock prediction and recommendation models. However, these models do not have deep learning methods as well as multi-long-term forecasting. A deep learning-based stock recommendation system was introduced in [19]. The technique improved a personalized stock suggestion but failed to use technical and fundamental stocks in a conceptualized form. A deep learning model-based stock investment recommendation system

was introduced in [20]. The system was highly performing but was black-boxed and failed to bring out definite investment insights.

Research Gap Identified: it is apparent that Transformer models [2-4] are also price forecast-specific (but not decision support). All recommendations systems [5-9] are signal based, rather than multi-horizon deep. Shallow learning is used by hybrid technical-fundamental models [10-12]. Unless interpretable, deep learning systems [6, 13, 19, 20]. The apparent necessity of an explainable, multi-horizon, hybrid forecasting and recommendation model necessitates the creation of a Temporal Fusion Transformer-based stock recommendation system provided in this.

In spite of the advancements in deep learning and financial analysis, there is still no system capable of performing all of the following together: making predictions for stock prices for various time periods, processing both technical and fundamental financial information, and providing adequate recommendations for buy, hold, and sell. As a result of this issue, the target of the current research is to design an easy and interpretable stock prediction and recommendation system by utilizing the Temporal Fusion Transformer. The proposed system should manage the deficiencies of previous systems by providing more accurate and more reliable investment advice in actual market environments.

3 METHODOLOGY

Stock price prediction, which is powered by deep learning, can be used to improve an Investment Recommendation System in the case of corporate finance. This paper presents a model that is built on the Temporal Fusion Transformer (TFT), which is a powerful deep learning architecture that is designed to forecast multi-horizon time series. The TFT model is a combination of attention and recurrent mechanisms to detect both the short-term variation and long-term trends in stock information. It is also capable of taking different classes of input features including known variables, observed variables, and static variables and is simple to comprehend. In this section, the methodology of the system development (data collection, preprocessing, feature engineering, model architecture, training and evaluation) is detailed. Data Collection in this study, there are two kinds of data that are used to examine the performance of stock; fundamental data and technical data. Yahoo Finance was used to retrieve the basic information of the chosen NIFTY 50 stocks over the last five years with the help of the finance library. The data has approximately 100 observations and contains

approximately 15 financial variables such as Net Income, Total Revenue, Operating Income, EBIT, Diluted EPS, Total Assets and Liabilities. These characteristics offer a regional review of the profitability of a company, financial condition and valuation of the company, which is essential in making long-term investment decisions. In the case of technical analysis, the data had been produced with the help of the ta library, and there were more than a thousand entries which had approximately 30 to 40 technical indicators. They can be the usual signs of SMA, EMA, MACD, RSI, Bollinger Bands, ATR and Aroon Oscillator, as well as volume indicators (On-Balance Volume (OBV) and Chaikin Money Flow (CMF)). These indicators are used together to determine the price trends, momentum and market volatility which assists in making better short-term trading and investment decisions. The current piece of work basically tries to create an elucidable and imprecise deep learning-based prescribing and anticipation of stock, which is a blend of technical and basic financial data trained by the Temporal Fusion Transformer (TFT). These are the following objectives:

O1. Gathering and pre-treating multi-source information spanning both past price information to technical indicators and core financial variables on the chosen NIFTY-50 stocks.

O2. The design of a hybrid forecasting framework capable of capturing both short-run and long-run temporal relationships in the prices of the stocks will be done using the Temporal Fusion Transformer.

O3. This paper will combine both an application of a technical indicator and a basic financial ratio into a single and comprehensive deep learning model to more effectively predict stock prices.

O4. This will be done by adding interpretability into the model by the variable selection and attention mechanisms of TFT to specify which variables are the most responsible in causing a change in stock price.

O5. Develop a recommendation system where the projected price trends and financial indicators are used to provide a signal of buying, holding, or selling.

O6. The performance appraisal of the proposed system could be examined with respect to the

MAPE and consistency of the recommendations of different stocks.

3.1 Data Pre-processing

Both the fundamental and technical datasets have undergone several preprocessing measures in order to be prepared to be modeled and predicted with some accuracy. In the case of the basic dataset, the missing of the features in the operating income, EBIT, Gross Profit and Total Expenses were filled in with statistical imputation methods like mean or median substitution based on the distribution of each feature. In other instances, when the missing data was small and not meaningful, then the concerned records were eliminated. The technical dataset was fully covered and it did not need imputation. After that, the feature engineering step was carried out to enhance the predictive power of the dataset. Numerical attributes, Open, High, Low, Close and Volume, were normalized using standard scaling so that they all have an equal contribution in the training of the models. Also, lag components were developed to make the model learn about past trends and time related trends that are imperative in prediction.

The information on time was also maintained by transforming date-time values into a numerical time index so that sequential models could inherently understand time based development. The label encoding was used to put the categorical features such as the stock ticker symbols through numerical encoding to fit the deep learning models. A chronological time-based split was then made between the data that was to be used as training (80 percent) and the data that was to be used as testing (20 percent). This method acts to preserve the time sequence of the events and does not allow future data to influence past model training as this is essential to avoid data leakage. As the Temporal Fusion Transformer (TFT) assumes time-series data to be in a structured format, the resulting dataset was reorchestrated into a (batch size, time steps, features) format. This format helps the model to identify both short term and long term dependencies and therefore, it can be used to predict stock prices as well as make credible investment recommendations.

3.2 Model Implementation

The Temporal Fusion Transformer (TFT) is a deep learning framework that is used as a stock recommendation system, but is a multi-horizon time series prediction architecture based on interpretable predictions. The TFT model combines both the fixed

and moving characteristics such as the technical indicators and the basic financial ratios, the dynamic characteristics, to identify intricate temporal relationships, uncertainty and the relevance of features over time, as well as giving effective and interpretable stock forecasts.

3.2.1 Variable selection networks (VSNs)

VSNs are applied to static covariates, past inputs, and known future inputs to dynamically select the most relevant features during training.

This allows the model to ignore irrelevant or noisy variables and focus only on impactful signals at each time step:

$$\alpha_i = \frac{e^{W_i x_i}}{\sum_j e^{W_j x_j}} \quad (1)$$

Where:

α_i is the attention weight,

W_i are learnable weights,

x_i is the input feature.

Gated Residual Networks (GRNs)

GRNs transform inputs while preserving stability and enabling non-linear combinations of variables. They apply gating mechanisms to regulate information flow and skip connections to ensure gradient stability.

The GRN operation is defined as:

$$\text{GRN}(x) = \text{LayerNorm} \left(x + \text{Gate}(W_2 \cdot \text{ELU}(W_1 x)) \right) \quad (2)$$

Where:

- W_1, W_2 are weights,

-ELU is the Exponential Linear Unit activation,

-LayerNorm ensures stable training,

- Gate is a sigmoid layer that controls feature propagation.

LSTM-Based Local Encoder

A recurrent encoder-decoder with LSTM cells is used to model local temporal dependencies and time-series dynamics.

It summarizes input sequences by capturing short-term memory, essential for detecting recent patterns in stock movements.

3.3 Self-Attention Layer

The transformer-based attention mechanism captures long-term dependencies across time steps, improving multi-step forecasting.

Multiple attention heads allow the model to weigh different time contexts in parallel.

The scaled dot-product attention is calculated as:

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V \quad (3)$$

Where:

Q, K, V are query, key, and value vectors, d_k is the key dimension, softmax ensures attention weights sum to 1

3.3.1 Quantile Regression Output

Instead of predicting a single value, TFT outputs multiple quantiles (e.g., 10th, 50th, 90th) to represent the range of possible future outcomes, supporting risk-aware recommendations.

The quantile loss is given by:

$$\mathcal{L}_q = \max(q(y - \hat{y}), (q - 1)(y - \hat{y})) \quad (4)$$

Where:

y is the true value,

\hat{y} is the predicted quantile value

q is the quantile (e.g., 0.5 for median).

Model Training

Dataset Split:

80% Training Set: Includes historical technical indicators and yearly fundamental data.

20% Testing Set: Contains unseen stock data from the final year to evaluate out-of-sample performance.

Loss Function

The model is trained using Quantile Loss, which is better suited for forecasting under uncertainty than Mean Squared Error.

It aids in the forecasting of not only a central value but a complete distribution of possible outcomes. The model adopts the AdamW optimizer that incorporates adaptive weight decay that enhances effective convergence so as to keep the model regularized. Our learning rate was 0.001, which is between training speed and stability. This enables the model to learn well without going too far and missing the best possible solutions. A batch of 64 samples per step provided enough memory to achieve constant learning rates. The training was not exceeded in terms of epochs, but the early-stopping was also in place to interrupt the training, when there was no further improvement. This is used to minimize the chances of overfitting. We also used a dropout rate of 0.3, which randomly kills neuron in the process of training. This

enhances the generalization capacity of the model as well as its strength with hidden data.

3.4 Hyper parameter tuning

We have conducted hyper parameter tuning through grid search to determine the optimal hyper parameters of Temporal Fusion Transformer (TFT) model. We experimented with hidden layer sizes of 16, 32, 64, and discovered that 64 gave the best results in terms of the representational capacity of the model in Gated Residual Networks (GRNs) and attention layers. We tested the attention head number at {1, 4, 8}, where 4 gave the most suitable trade-off between complexity and capturing long term dependencies. We experimented with learning rates of {0.01, 0.001, 0.0001} and chose 0.001 which we found to be stable and efficient. We tried dropout rates of {0.2, 0.3, 0.5} and found that 0.3 works well in preventing overfitting as well as not impeding learning to an extent. We tested batch sizes of: 32, 64, 128 and 64 gave a good balance between the computational and the generalization of the model. We have also chosen the context length which is the number of days on which the model depends on the historical data (we have tried 30 and 60 days as well as 90 days) and found 60 days to be the best. Lastly, we maximized the forecast horizon or the number of days the model predicts and it was 14 days. This makes it possible to have the correct mid-term forecast which is appropriate to make investment decisions.

3.4.1 Predictions and model evaluations

The Temporal Fusion Transformer (TFT) model predicts stock closing prices over a 30-day forecast horizon. This helps inform investment decisions based on both historical trends and financial fundamentals. The predicted prices serve as the basis for a recommendation engine that combines technical signals and key financial ratios to classify stocks into buy, hold or sell categories.

To evaluate the accuracy of these predictions, the Mean Absolute Percentage Error (MAPE) was used as the primary performance metric. MAPE quantifies the average deviation between predicted and actual prices and is defined as:

$$MAPE = \frac{1}{n} \sum_{i=1}^n \left| \frac{A_i - P_i}{A_i} \right| \times 100 \quad (5)$$

where:

A_i represents the actual closing price,

P_i represents the predicted closing price
 n is the total number of test samples

The smaller the MAPE value, the greater the model accuracy and the prediction ability. The forecasted closing prices within the 30 days forecast period were checked against real values of the test data. This comparison graphically proved the model fitting with market changes and trends. It was demonstrated that the TFT model was able to describe not only the short-term trends but also the long-term trends, which are crucial in transforming the financial situation. The results of the prediction model were forwarded to the recommendation engine. Other important indicators that were incorporated in this engine were Earnings per Share (EPS), Price-to-Earnings ratio (P/E), Return on Equity (ROE), Debt to Equity ratio, Net Income, and Total assets. These elements provide a good image of the real worth of a company and its financial welfare. These financial fundamentals were compared with the predicted trends and analyzed using the engine to generate recommendations. This will help in making sure that the decisions being taken are not merely reactive to price fluctuations but they are also pegged on the financial strength of a stock. A structured workflow (Figure 2) was used to demonstrate the whole process, including the preparation of data and prediction using TFT as well as generating recommendations. It begins with the data collection (technical, fundamental), preprocessing and sequence generation. Training and evaluation of the TFT model is then done using MAPE. Lastly, forecasted values/financial wisdom are synthesized to provide high quality and data-driven investment recommendations.

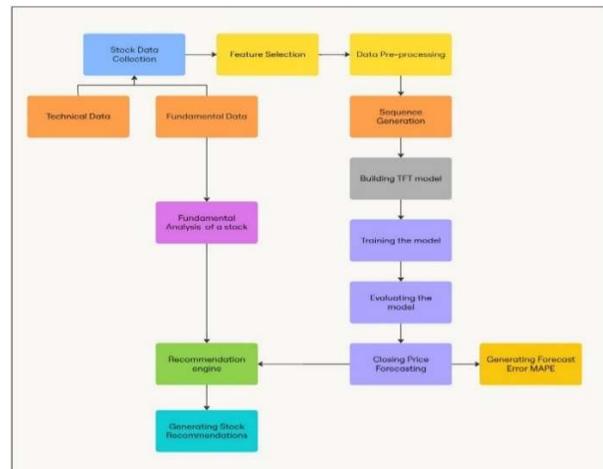


Figure 2: Proposed model

The Figure 2 shows an entire workflow of stock price prediction and recommendation with the help of

Temporal Fusion Transformer (TFT) model. It starts with the collection of stock data that involves the collection of both technical (including stock price indicators and volume-based measures) and fundamental (including financial statements and market ratios) data. Once data is gathered, a feature selection process is undertaken in order to reduce the data to the most pertinent predictors. The data is then refined and data preprocessing is implemented after which missing values, scaling and formatting modifications are carried out to model the data. After preprocessing, the data is then transformed into sequences, which are applicable in deep learning, by a sequence generation step. These sequences are subsequently made to construct the TFT model, followed by training of the model with the use of past stock data. The model is tested after the training to determine the performance and accuracy of the model. It generates the predictions in the form of closing stock price forecasts and the accuracy of the forecast is measured with the help of Mean Absolute Percentage Error (MAPE). In parallel to the forecasting pipeline, a fundamental analysis of the stock using the collected technical and basic data is also conducted to determine the potential of an intrinsic value and the stock performance in the long term. Both forecasting and fundamental analysis outcomes are input into a recommendation engine, which combines predictive information and financial data to produce investment-related recommendations. Last but not least the system will provide stock recommendations, which will help the investors make sound financial decisions.

4 RESULT AND DISCUSSION

The Temporal Fusion Transformer (TFT) model was found to be optimistic in its outcomes to predict stock prices and give an investment recommendation. It was able to predict the closing prices of the stock within 30 days using past historical stock patterns, basic values as well as dynamic technical values. The primary performance indicator was the Mean Absolute Percentage Error (MAPE) and remained at a relatively low level in all the stock. This was a sign of high generalization and accuracy of prediction. The time-series comparison of actual and predicted prices revealed that the TFT model could best describe time-related complex patterns particularly when the markets were volatile. Its multi-horizon forecasting was the ability to

recognize the short-term and long-term trends more effectively. This was reflected to the fact that prediction curves were smoother in stable periods and responsive in market fluctuations [21,22]. The TFT-based system in the recommendation module allowed the categorization of stock based on the predicted price trends and the annual fundamental ratios as buy, hold, or sell. The combination of these stationary and time-dependent characteristics made it possible to consider the potential to invest more thoroughly. The stocks that are likely to increase and have good financial position were indicated as something to buy and the ones that were indicating a weakening position were indicated as something to sell or retain. The ability to use visuals comparing the movement of predicted and actual stock prices was also an indicator of the validity and simplicity of the model. The TFT model mechanisms of attention helped to shed light on the significance of the most influential variables in the predictions to enhance the clarity of investment decisions. These findings affirm the utility of the TFT to the short and long-term trading and portfolio management and therefore qualifies as a useful instrument in the current data-driven investment world.

4.1 RELIANCE

The TFT-based model predicted Reliance's stock price with a Mean Absolute Percentage Error (MAPE) of 3.7%, showing moderate prediction accuracy. The expected downward trend led to a "sell" recommendation. While short-term technical indicators showed weakness, broader financial stability and long-term positioning stayed strong. This case highlights the model's ability to signal short-term risks, providing timely warnings to investors.

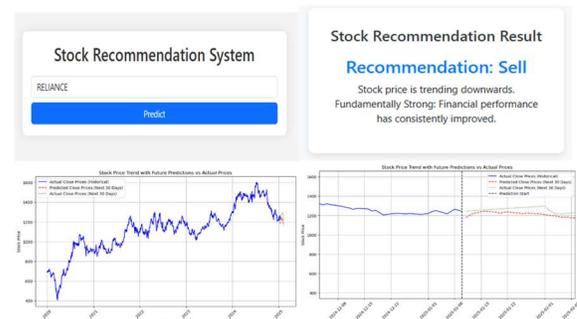


Figure 3: Stock Recommendation System

4.2 HINDUNILVR

The model showed a MAPE of 2.75%, indicating high accuracy in predicting Hindustan Unilever's short-term price changes. The forecasted decline led to a "sell" recommendation. Although the company

has strong long-term fundamentals, recent market trends suggested short-term corrections. This demonstrates that even fundamentally strong companies can face temporary setbacks, highlighting the importance of making tactical adjustments.

performance to highlight promising investment opportunities.

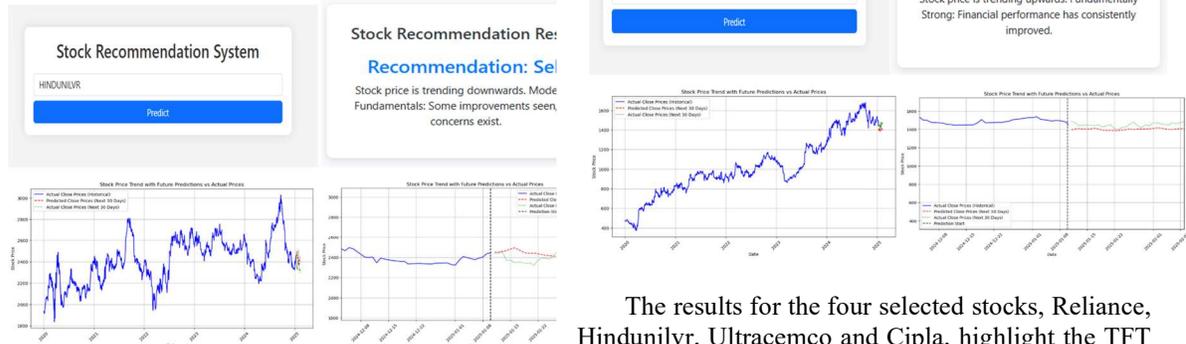


Figure 4: Stock Recommendation System

4.3 ULTRACEMCO

For Ultratech Cement, the model produced a MAPE of 6.77%. This indicates greater price volatility and prediction complexity. The recommendation engine suggested a "hold", backed by slight upward movement in the market. Although technical indicators were unclear, stable financial data supported a careful approach. This situation highlights the need for balanced judgment in unpredictable stock markets.

The results for the four selected stocks, Reliance, Hindunilvr, Ultracemco and Cipla, highlight the TFT model's ability to predict short-term price trends with significant accuracy. It also generates recommendations that show real market behavior.

Table 1: MAPE Values and Recommendations of Selected Stocks

| Stocks | MAPE | Actual | Predicted |
|------------|-------|--------|-----------|
| Reliance | 3.70% | Sell | Sell |
| Hindunilvr | 2.75% | Sell | Sell |
| Ultracemco | 6.77% | Hold | Hold |
| Cipla | 3.29% | Buy | Buy |

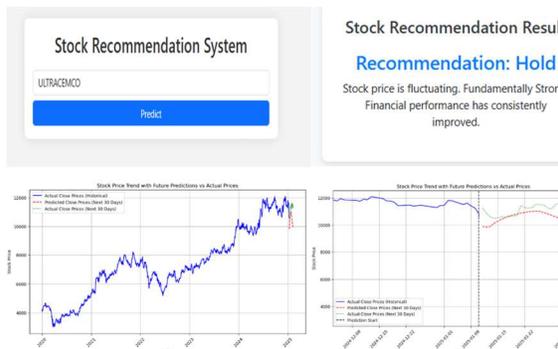


Figure 5: Stock Recommendation System

4.4 CIPLA

Cipla's forecast reached a MAPE of 3.29%. This shows good prediction accuracy and indicates a moderate upward trend. The model suggested a "buy," matching real market behavior and backed by strong financial indicators. This case demonstrates how the TFT model effectively combines trend recognition with company

Table I gives an overview of the performance of TFT model. The smallest MAPE was of Hindunilvr with 2.75% and Cipla with 3.29%. This indicates high accuracy of forecasting. Recommendations of the model were close to the real movement in the market. The observed declines were associated with the sells signals of Reliance and Hindunilvr. There was positive momentum in Cipla as shown by the buy signal. In the case of Ultracemco, the model was able to address the uncertain market condition despite the relatively high MAPE of the input of the hold recommendation. On the whole, the findings prove that the TFT model is reliable in integrating technical forecasting and financial information to make investment decisions.

In this way, the majority of the current stock prediction and recommendation systems replicate one of the two approaches: technical indicators [5 -9] or financial fundamentals [11,12], yet very few of them is based on a combination of both in a deep learning system. The hybrid models, which combine these two sources, are always better than the single source models as demonstrated by Nti et al. [10] since they are effective in capturing both the market behavior of a company as well as the value of the company itself.

Nevertheless, most of the previous hybrid systems have employed shallow neural networks or basic rule of thumb approaches, which cannot capture long-term and intricate temporal interactions in stock market data. Conversely, this paper uses Temporal Fusion Transformer (TFT), which is purposefully created to consume various forms of financial inputs, such as static company fundamentals, time-varying market prices and technical indicators, and known future inputs [24]. This renders TFT very appropriate in dealing with mixed financial information in a way that is organized and understandable. In contrast to the traditional deep learning models [6,13,19,20], TFT employs Variable Selection Networks (VSN) to learn the significance of all the technical and fundamental features at each time step automatically [2]. Consequently, less relevant or noisy indicators are not regarded as important and more powerful factors are considered and the transparency of the model is maintained. Thus, the process of feature selection in this context is not random, but it is trained directly on the data, which enhances the predictive quality as well as the quality of investment decisions [24].

Table 2. Comparison of Proposed Work with Prior Literature

| S.No | Prior Literature | Limitation in Prior Work | Proposed Work |
|------|------------------|------------------------------------------------------|---------------------------------------------------------------------------------|
| 1 | [1,6,13] | Cannot model complex long-term temporal dependencies | Uses Temporal Fusion Transformer (TFT) for multi-horizon deep temporal modeling |
| 2 | [2,3,4] | No integration with investment recommendation | TFT forecasts are directly used for buy/hold/sell decisions |
| 3 | [5,9] | Partial view of market and company performance | Integrates technical, fundamental, and price data in a single model |
| 4 | [10,12] | Cannot learn dynamic feature importance | TFT uses Variable Selection Networks for dynamic feature weighting |

Table 2 shows that Most stock prediction models currently available in the stock markets have some deficiencies. Conventional deep learning models such as LSTM and CNN primarily aim at predicting stock prices in the short term. They fail when it comes to analyzing stock markets in terms of their overall trends in the long term as well as changes in their behavior in the

short term. Due to this fact, these models fail in stock investment for a medium term as well as a long term. However, the proposed system is built using the Temporal Fusion Transformer (TFT) technology, which is generally capable of predicting stock prices in multiple time terms.

Recent examples of transformer networks applied to problems with the TFT model of stock prices have shown that the model has a better performance at stock price forecasting. Nonetheless, there is no support system within these works, which helps to make decisions or suggests buying, holding, or selling stocks. In our project, an extension of the TFT model with a support system for buying, holding, or selling decisions based upon stock prices determined through forecasting would allow us to utilize the model more effectively.

5 CONCLUSION

The rationale behind this study is the necessity to have a stock prediction system that is accurate as well as interpretable and really close to the way investment decisions are made in real life. Most of the current stock prediction methods are based on technical indicators and price forecasting, or black box deep learning models that are not transparent. These restrictions decrease their practical usefulness particularly in volatile markets where the investors need detailed and accurate guidance. To deal with these issues, this study presents a hybrid deep learning structure developed based on the Temporal Fusion Transformer (TFT). The model incorporates the technical indicators, past prices data and both basic financial variables into a single multi-horizon forecast system. The model enhances the quality of a prediction as well as offering interpretability since it identifies which factors contribute more significantly to a given prediction using a combination of variable selection networks and attention mechanisms. Moreover, a recommendation engine converts a forecast of price changes and financial well-being of the company into an actual buy, hold or sell, in effect closing the gap between forecasts and investment choices. The experimental analysis of the sampled NIFTY-50 stocks shows minimal prediction errors and investment suggestions that are way much related to the actual market behavior. These findings support the validity and power of the integration of technical and fundamental data into a sensible learning model which can be more dependable than traditional or single-source models.

Nevertheless, the study has its weaknesses. It is based on a short historical window and small sample

of stocks and fails to include external factors like news sentiment, macroeconomic factors and geopolitical factors. As much as TFT enhances interpretability, analysis of attention outputs cannot be completed effectively without domain expertise. The suggested framework will be a great addition to the currently available literature as it combines deep learning-based forecasting with understandable, finance-oriented decision support. Its applicability in practice-oriented investment is high, and the following research can improve the performance by the use of sentiment analysis, macroeconomic factors, reinforcement learning and the use of risk-adjusted strategies in dynamic financial settings.

This paper makes the following significant contributions to the field of stock market forecasting and investment recommendation:

Contributor 1: Implementation of A unified deep learning framework is proposed that integrates historical stock prices, technical indicators, and fundamental financial data within a single Temporal Fusion Transformer (TFT) architecture, enabling holistic modelling of stock market behaviour.

Contributor 2: Checking how to the TFT provides not only accurate multi-horizon stock price forecasts but also built-in interpretability through variable selection networks and attention mechanisms, allowing investors to understand which factors influence predictions.

Contributor 3: validated using NIFTY-50 stocks, demonstrating low MAPE values and strong alignment between predicted trends and real market movements.

Contributor 4 and 5 : Paper write up and arranging in the journal format

REFERENCES

- [1]. Ariyo, Adebisi A., Adewumi O. Adewumi, and Charles K. Ayo. "Stock price prediction using the ARIMA model." 2014 UKSim-AMSS 16th International Conference on Computer Modelling and Simulation, pp. 106–112. IEEE, 2014.
- [2] Lim, Bryan, Sercan Ö. Arık, Nicolas Loeff, and Tomas Pfister. "Temporal fusion transformers for interpretable multi-horizon time series forecasting." *International Journal of Forecasting* 37, no. 4 (2021): 1748–1764.
- [3] Hu, Xiaokang. "Stock price prediction based on temporal fusion transformer." 2021 3rd International Conference on Machine Learning, Big Data and Business Intelligence (MLBDBI), pp. 60–66. IEEE, 2021.
- [4] Ho, Raymond, and Kevin Hung. "Ceemd-based multivariate financial time series forecasting using a temporal fusion transformer." 2024 IEEE 14th Symposium on Computer Applications & Industrial Electronics (ISCAIE), pp. 209–215. IEEE, 2024.
- [5] Lokesh, D., V. Vani, and N. Karthik. "Stock Recommendation System for Better Investment Plan." 2024 International Conference on Signal Processing, Computation, Electronics, Power and Telecommunication (IConSCEPT), pp. 1–6. IEEE, 2024.
- [6] Li, Jianlong, Siyuan Wang, Zhihang Zhu, Minghao Liu, Changjiang Zhang, and Bingyan Han. "Stock Prediction Based on Deep Learning and its Application in Pairs Trading." 2022 International Symposium on Networks, Computers and Communications (ISNCC), pp. 1–7. IEEE, 2022.
- [7] Sundarakamatchi, S., and M. S. Gajanand. "Recommendation Engine for Stock Market Trading." In *Computational Management: Applications of Computational Intelligence in Business Management*, pp. 499–516. Springer, 2021.
- [8] Hernández-Nieves, Elena, Álvaro Bartolomé del Canto, Pablo Chamoso-Santos, Fernando de la Prieta-Pintado, and Juan M. Corchado-Rodríguez. "A machine learning platform for stock investment recommendation systems." In *Distributed Computing and Artificial Intelligence*, 17th International Conference, pp. 303–313. Springer, 2021.
- [9] Vismayaa, V., K. R. Pooja, A. Alekhya, C. N. Malavika, Binoy B. Nair, and P. N. Kumar. "Classifier based stock trading recommender systems for Indian stocks: An empirical evaluation." *Computational Economics* 55 (2020): 901–923.
- [10] Nti, Isaac Kofi, Adebayo Felix Adekoya, and Benjamin Asubam Weyori. "A systematic review of fundamental and technical analysis of stock market predictions." *Artificial Intelligence Review* 53, no. 4 (2020): 3007–3057.
- [11] Namdari, Alireza, and Zhaojun Steven Li. "Integrating fundamental and technical analysis of stock market through multi-layer perceptron." 2018 IEEE Technology and Engineering Management Conference (TEMSCON), pp. 1–6. IEEE, 2018.

- [12] Beyaz, Erhan, Firat Tekiner, Xiao-jun Zeng, and John Keane. "Comparing technical and fundamental indicators in stock price forecasting." 2018 IEEE 20th International Conference on High Performance Computing and Communications; IEEE 16th International Conference on Smart City; IEEE 4th International Conference on Data Science and Systems (HPCC/SmartCity/DSS), pp. 1607–1613. IEEE, 2018.
- [13] Wang, Weina, and Krishn Kumar Mishra. "A novel stock trading prediction and recommendation system." *Multimedia Tools and Applications* 77, no. 4 (2018): 4203–4215.
- [14] Bag, Vipul, and Uday V. Kulkarni. "Stock price trend prediction and recommendation using cognitive process." *International Journal of Rough Sets and Data Analysis (IJRSDA)* 4, no. 2 (2017): 36–48.
- [15] Paranjape-Voditel, Preeti, and Umesh Deshpande. "A stock market portfolio recommender system based on association rule mining." *Applied Soft Computing* 13, no. 2 (2013): 1055–1063.
- [16] Tirea, Monica, Ioan Tandau, and Viorel Negru. "Stock market multi-agent recommendation system based on the Elliott wave principle." In *Multidisciplinary Research and Practice for Information Systems*, pp. 332–346. Springer, 2012.
- [17] Hegde, Manoj S., Ganesh Krishna, and Ramamoorthy Srinath. "An ensemble stock predictor and recommender system." 2018 International Conference on Advances in Computing, Communications and Informatics (ICACCI), pp. 1981–1985. IEEE, 2018.
- [18] Sharma, Vikram, Manik Rakhra, and Gauri Mathur. "*Hybrid approaches for stocks prediction and recommendation system.*" In *E3S Web of Conferences*, vol. 453, p. 01047. EDP Sciences, 2023.
- [19] Wijerathne, Narada, Jamini Samarathunge, Krishalika Rathnayake, Supuni Jayasinghe, Sapumal Ahangama, Indika Perera, Vinura Dhananjaya, and Lushanthan Sivaneasharajah. "*Deep learning based personalized stock recommender system.*" In *International Conference on Neural Information Processing*, pp. 362–374. Springer, 2023.
- [20] Parzyszek, Michal, and Stanislaw Osowski. "Deep Learning Recommendation System for Stock Market Investments." In *International Work-Conference on Artificial Neural Networks*, pp. 263–275. Springer, 2023.
- [21] D. A. Tepper et al., "Multi-Sensor Temporal Fusion Transformer for Stock Performance Prediction: An Adaptive Sharpe Ratio Approach," *Sensors**, vol. 25, no. 3, p. 976, Feb. 2025, doi: 10.3390/s25030976.
- [22] S. Selim Çalık, A. Akyüz, Z. H. Kilimci, and K. Colak, "Explainable-AI powered stock price prediction using time series transformers: A Case Study on BIST100," *NEP RePEc**, Jun. 2025, doi: (arXiv preprint)