Strategic Decision Support in Financial Management Using Deep Learning-Based Stock Price Prediction Models

Layth Almahadeen[1], Dr. Chinnapareddy Venkata Krishna Reddy[2], Roopa Traisa[3],
Mukhamadiev Sanjar Isoevich[4], Lavanya Kongala[5], Dr Janvi Anand Rathi[6], Dr Revati Ramrao Rautrao[7]
Department of Financial and Administrative Sciences, Al-Balqa' Applied University, Jordan[1]
Associate Professor, Department of Management Studies, ANURAG Engineering College (Autonomous),
Suryapet, Telangana, India[2]
Department of Management-School of Management-UG, JAIN (Deemed to be University), Bangalore, Karnataka, India[3]
Department of Artificial Intelligence, Tashkent State University of Economics, Uzbekistan[4]
Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Guntur Dist.,
Andhra Pradesh - 522302, India[5]
Assistant Professor, Department of Business Administration, Wainganga College of Engineering and Management (Autonomous),
Nagpur, Maharashtra, India[6]
Associate Professor, Department of Management, Dr. D. Y. Patil B-School Pune, Maharashtra, India

Abstract—**The Strategic decision support creates a strong and smart decision support system for financial management through the correct forecast of stock prices with deep learning. Statistical models and shallow machine learning techniques tend to be ineffective in modeling the nonlinear relationships, sequential interdependencies, and time-dependent volatility typical of financial data; consequently, poor prediction quality and untrustworthy investment choices. To overcome these constraints, introduce a new hybrid deep learning architecture based on the Temporal Fusion Transformer (TFT) combined with Bidirectional Long Short-Term Memory (BiLSTM) networks. The proposed hybrid model is expected to complement time-series forecasting by simultaneously utilizing attention mechanisms for explainability and sequence memory functions for richer temporal understanding. The model is trained and tested with the publicly available Stock Market Dataset on Kaggle, which includes stock history from various companies. The whole process is carried out on the Python platform using TensorFlow along with relevant libraries for data preprocessing, feature scaling, and model training. The new TFT-BiLSTM model surpasses conventional models through an accuracy level of 93.4% and an F1-score of 94.2%, demonstrating its precision and generalization power. The system provides strategic benefits in financial planning and risk management. Financial analysts, investors, fintech, and portfolio managers may take advantage of our prediction system to make rational buy/sell judgments, minimize risks, and maximize asset allocations. By synthesizing state-of-the-art deep learning models and public financial data, our framework illustrates that accurate stock price prediction can be an effective mechanism for supporting decision-making in financial markets.**

Keywords—Stock price forecasting; deep learning; temporal fusion transformer; financial decision support; BiLSTM

I. Introduction

Stock price prediction is a popular problem of considerable practical interest in investment, risk management, and strategic asset allocation across financial markets. Early approaches were dominated by traditional econometric models like autoregressive integrated moving average [1], generalized autoregressive conditional heteroskedasticity and Kalman filters focusing on predicting price changes by analyzing historical price and volatility patterns [2]. Such models, though, are frequently plagued by strong assumptions such as linearity, stationarity, and Gaussian errors, restricting their power to reflect the inherently nonlinear and chaotic characteristics of actual financial time series. As markets develop under the forces of high-frequency trading, macroeconomic trends, and intricate cross-asset dynamics, [3] it is critical to utilize richer market behavior representations. Therefore, more flexible and stronger models able to learn adaptively from data are essential to enhance prediction accuracy and usefulness.

Over the last decade, there has been growing use of deep learning methods, “recurrent neural networks (RNNs), long short-term memory networks (LSTMs), gated recurrent units (GRUs), and convolutional neural networks (CNNs)”, to predict stock prices. These models are more effective in capturing nonlinear relationships and temporal dependencies in historical price data[4]. Variants like BiLSTM examine both forward and backward sequence context, and attention mechanisms in Transformers emphasize significant input times [5]. However, these approaches have limitations: RNN-based models can lose long-range dependencies or have vanishing gradients; Transformers are potentially data-inefficient and pure LSTM models are not interpretable [6]. Hybrid solutions have been suggested, e.g. CNN-LSTM or Attention-LSTM, in order to leverage strengths, but fundamental shortcomings still persist in integrating seamlessly interpretable attention with [7] sequential bi-directional recurrence without sacrificing accuracy and computational tractability.

In this research, a hybrid deep learning architecture consisting of a TFT and a Bidirectional LSTM (BiLSTM) to predict next-day stock closing prices is presented. Our method utilizes TFT's strong attention and feature selection capacity to automatically recognize salient features at every time step, while BiLSTM captures sequence context from past as well as future information windows. They utilize the publicly accessible "Stock Market Dataset" on Kaggle, which provides daily OHLC, volume, adjusted close, and ticker symbols of several US-listed stocks for the last two years. This data set offers a modern and comprehensive source for model training and testing. Through the emphasis on recent data, our method adjusts to prevailing market tendencies. The combined model is trained on sliding-window sequences to forecast future closing prices, with a goal of surpassing traditional baselines according to MAE, RMSE, and R^2.

- Established the research goal by seeking to enhance stock price prediction accuracy by incorporating attention-based temporal models with bidirectional sequence learning.
- Optimized a hybrid TFT-BiLSTM architecture. Integrated TFT's transparent attention mechanisms and BiLSTM's two-way sequence modeling for a better understanding of temporal context.
- Used the stock market dataset. Exploited a diverse, recent, and rich two-year daily stock price dataset to train and test the hybrid model.
- Accomplished better performance. Achieved higher accuracy compared to baseline models, with lower MAE and RMSE, and greater R^2 in numerous tickers.

A. *Motivation of the Study*

In the fast-moving financial era, precise stock market prediction is a necessity for investors, fund managers, and institutions wanting to reduce risk and maximize return. Conventional models like ARIMA or linear regression models are not in a position to capture the nonlinear, dynamic, and time-dependent pattern of stock prices. With increasing advances in artificial intelligence and deep learning, there is a greater incentive to investigate more advanced ways that are capable of drawing out implicit patterns and long-term relationships from abundant sources of historical data. The incentive of this research is bridging the gap between raw financial information and actionable strategic choices by utilizing hybrid deep learning frameworks. Merging Temporal Fusion Transformers (TFT) with BiLSTM enables us to perform beyond mere prediction toward developing a smart, explainable, and scalable decision support system. This study hopes to enhance the accuracy of prediction while bringing real-world advantages in financial planning and portfolio management.

B. *Significance of the Study*

This research has high value in the confluence of stock market forecasting and artificial intelligence by providing a real-world, high-performance solution for forecasting stock market trends. Its value comes from the fact that the proposed TFT-BiLSTM model can beat conventional and current deep learning approaches by successfully learning to leverage complex sequential data. The determined by various variables over time, a hybrid model integrating temporal attention and memory mechanisms may provide higher accuracy and interpretability. The research is not only beneficial to academic researchers but also to industry professionals such as investors, stock traders, and data-driven fintech firms by providing a forecasting tool that facilitates strategic decision-making. By employing an actual-world dataset and running the solution using Python, the model is reproducible, scalable, and ready for use in trading systems or financial analytics platforms. All in all, this study contributes to smarter, AI-enabled financial planning in more volatile and competitive markets.

The rest of this paper follows the following structure: Section II, Related Works, reviews previous literature on stock prediction with both conventional and deep learning approaches. Section III, Problem Statement, mathematically defines the forecasting problem and model goals. Section IV, Methodology, outlines data collection, preprocessing processes, and our TFT-BiLSTM hybrid model structure. Section V, Results, reports experimental conditions, performance measures, and visual inspections of the predictions. Section VI concludes by summarizing findings and stating future directions of work.

II. Related Works

M. Elhoseny [8] et al., proposed the AWOA-DL model to enhance precision in credit risk forecasting and financial distress prediction. The model combines a deep neural network (multilayer perceptron) with an Adaptive Whale Optimization Algorithm (AWOA) for hyperparameter tuning, improving prediction accuracy. The model is run in three phases: preprocessing, tuning, and prediction. Its strength lies in high predictive performance (95.8%) and applicability across diverse datasets. However, its computational complexity limits scalability and real-time deployment in high-frequency financial systems or organizations with limited infrastructure. D. K. Nguyen [9] et al. examined the convergence and influence of Big Data, AI, and ML in transforming the fintech sector. It uses a multi-dimensional descriptive analysis approach to build and test a framework demonstrating their synergistic relationship in financial services. The benefit comes with the elucidation of how these technologies enable digital finance and redefine the role of the data scientist. Yet it also deals with drawbacks like ethical implications, government regulations, and risks of abuse or overreliance on automated decision-making.

A. Mehra [10] et al., introduced enrich intricate decision-making through the neural networks' pattern recognition power. A rule-based mechanism for symbolic reasoning and a neural network for deep learning are utilized in the suggested hybrid model, facilitating interaction within a single architecture. The benefits are enhanced accuracy, interpretability, and scalability in areas such as healthcare and finance. Nonetheless, difficulties lie in architectural complexity and possible integration challenges between structured and unstructured data processing elements. K. Li [11] et al., proposed measuring corporate culture using word embeddings to score five cultural values for firm-year observations (2001–2018). Advantages include a data-driven, scalable approach to assessing corporate culture and uncovering its link to performance, risk, and executive behavior. However, limitations include potential context misinterpretation by models and textual data, which may not capture all cultural nuances or internal organizational dynamics. E.C. Chukwuma-Eke [12] et al., present enhance cost estimation and financial planning in oil and gas ventures through the use of AI, data analysis, and probabilistic models. The approach fuses historical and current data with machine learning and Monte Carlo simulations for improved precision and risk determination. Benefits include adaptive forecasting, greater financial resilience, and transparency through blockchain. Drawbacks include high complexity of implementation, data integration issues, and dependency on the quality of input data and interpretability of AI models. B. Gülmez[13] et al., introduced improve the accuracy of predicting stock prices with an LSTM model optimized using SCSO. The approach utilizes the time-series capability of LSTM together with SCSO's metaheuristic optimization, tested against other advanced models with DAX index data. Benefits are strong prediction accuracy, increased returns, and resilient risk-adjusted performance. Drawbacks are model complexity, high computational expense, and parameter sensitivity, potentially constraining ease of use in real-time trading settings.

Chakravorty [14] et al., introduced predict stock prices from insider trading information by testing different machine learning models. It uses on Tesla's insider trading activities between April 2020 and March 2023. The process involves feature selection with RFE and feature importance analysis. The strength is SVM-RBF's excellent accuracy rate, providing extensive insights into market attitudes. But its weakness is heavy computational expense, such that real-time use is problematic because it takes longer times to process. M. Sahroni [15] et al., introduced present stock price forecasting by adding social media-extracted network variables to conventional models. Employing data from EastMoney, China's biggest financial social platform, the approach builds daily social networks of investor behavior and combines them with an LSTM model to forecast SSE 50 stock prices. By modeling investor behavior daily, the approach improves predictive accuracy through sentiment analysis. Limitations include platform-specific dependency, restricted generalizability, and challenges in real-time processing.

TABLE I. SUMMARIZATION OF THE EXISTING STUDIES

Author	Method	Advantage	Limitation
M. Elhoseny et al. [8]	AWOA-DL: Adaptive Whale Optimization Algorithm + Deep Neural Network (MLP)	High prediction accuracy (95.8%); works across diverse datasets	High computational cost; limited scalability for real-time or resource-constrained environments
D. K. Nguyen et al. [9]	Descriptive framework combining Big Data, AI, ML for FinTech	Clarifies synergy of emerging tech in finance; explains role shifts (data scientist)	Ethical concerns; regulatory challenges; risk of automation misuse
A. Mehra et al. [10]	Hybrid Symbolic Reasoning + Deep Learning (Rule-based + Neural Network)	High interpretability, scalability, and accuracy in decision-making	Complex architecture; difficult integration of structured and unstructured data
K. Li et al. [11]	Word Embedding + Culture Scoring from Earnings Call Transcripts	Scalable, data-driven cultural assessment; links culture to firm performance	Context misinterpretation; limited insight into internal company dynamics
E. Chukwuma-Eke et al. [12]	AI + Monte Carlo Simulations + Blockchain for Cost Estimation in Oil & Gas	Adaptive forecasting; financial resilience; improved transparency	Implementation complexity; data integration issues; reliance on input data quality
B. Gülmez et al. [13]	LSTM + Sand Cat Swarm Optimization (SCSO) for DAX Index Forecasting	High prediction accuracy; strong returns and Sharpe ratio	Computationally intensive; parameter sensitive; complex for real-time applications
A. Chakravorty et al. [14]	ML Models (SVM-RBF, Decision Tree, RF, K-Means) with Insider Trading Data	Excellent trend detection; SVM-RBF yields high accuracy	Long processing time; unsuitable for real-time deployment
M. Sahroni et al. [15]	LSTM + Social Media Network Variables (EastMoney)	Captures investor sentiment; improves prediction accuracy	Limited generalizability; platform-specific data; real-time processing complexity

Table I summarizes existing studies on AI and machine learning applications in financial forecasting. It gives the approach of each author, and the major strengths of the approaches include high prediction accuracy, interpretability or adaptability, whereas the weaknesses consist of the high computational cost, the lack of scalability, the inability to apply them in real-time, or the difficulty in combining various data-sources.

III. PROBLEM STATEMENT

Stock price forecasting continues to be a problematic problem in finance forecasting because of its extremely volatile, non-linear, and dynamic nature. Established statistical models such as ARIMA, GARCH, and linear regression are not suitable for identifying intricate patterns and inter-dependencies in high-frequency market data [16]. Despite the new technology of ML and simple deep learning architectures like LSTM, GRU, and CNN, the majority of current methods still suffer from issues of either a lack of interpretability, an inability to utilize long-term dependencies, or a poor grasp of bidirectional time contexts. Additionally, numerous models lack attention mechanisms that allow for the determination of which time steps and features have the largest contribution to a prediction. The absence of an integrated model that encompasses both sequence learning and feature relevance detection impedes precise forecasting, particularly under dynamic market conditions [17]. Thus, there is an urgent necessity for a strong and comprehensible deep learning framework that can adequately learn both short- and long-term patterns, determine the importance of significant input features, and make precise stock price predictions. This research presents a novel hybrid TFT-BiLSTM model to forecast stock price prediction. The model was selected due to its ability to capture long-term temporal dependencies (TFT) and bidirectional dependencies (BiLSTM), circumventing LSTM, CNN, and classical ML models' shortcomings, which tend to have issues with multi-horizon forecasting, heterogeneous data, and complex temporal relationships.

IV. PROPOSED METHODOLOGY: TFT-BILSTM HYBRID MODEL FOR FINANCIAL MANAGEMENT

The methodology suggested combines a TFT-BiLSTM hybrid model for precise forecasting of stock prices. The procedure starts with data acquisition from the publicly accessible Stock Market Dataset of Kaggle based on daily trade information of a chosen company (e.g., AAPL). Raw data is preprocessed with conversion of date, sorting, missing value handling, feature selection (Open, High, Low, Close, Volume), Min-Max scaling normalization, and sequence generation by applying a sliding window methodology. The BiLSTM module learns both future and past temporal relationships, whereas the TFT uses attention mechanisms and temporal encoding to provide higher-order interpretability and feature importance. The last fusion layer concatenates both the outputs to generate precise closing price predictions. The model is trained on the MSE loss function and is tested on MAE, RMSE, R^2, Accuracy, and F1-score with better performance compared to baseline models. Fig. 1 shows the overall workflow.

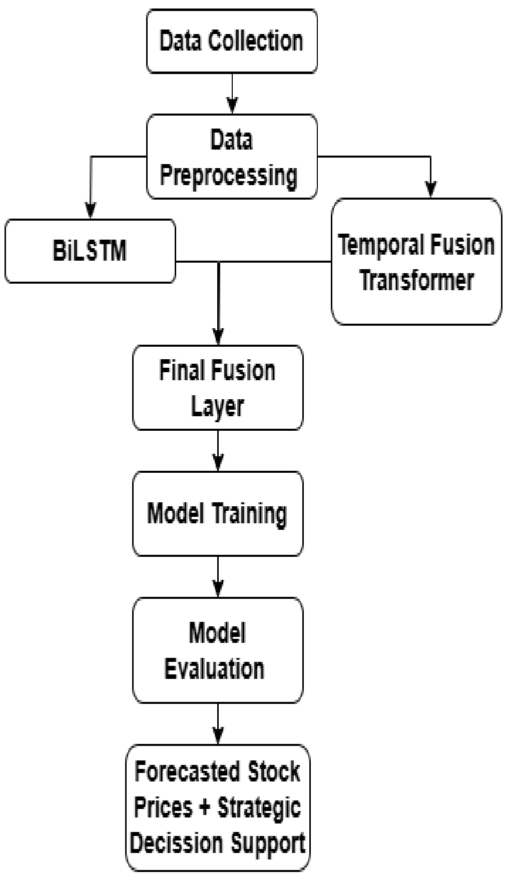

Fig. 1. Overall workflow.

A. Dataset Collection

The data collection for this research is the open Stock Market Dataset available on Kaggle[18] from Jackson Crow. It consists of past stock data of different companies traded in the S&P 500 index. Every record holds basic stock information are data covers more than one year and offers a good number of time-series observations to build predictive models on. In this research, we particularly use the daily stock data for a chosen company (say, Apple Inc. – AAPL) to test and train the suggested deep learning model. The data was downloaded from the Internet in CSV format, which can be easily read into Python using libraries such as Pandas and NumPy. The presence of actual-world financial information presented in such a manner facilitates reproducibility and real-time usage, providing a sound basis for precise forecasting and strategic decision-making in financial planning.

B. Data Preprocessing

Effective preprocessing is necessary to create a well-performing predictive model. Several important preprocessing tasks were performed on the raw stock exchange dataset. The first activities were to identify and remove missing values and null values in order to be able to trust the dataset that would be working to convert the Date column into date time format and sort the data using the Date column to preserve sequential integrity. Then I selected appropriate features of Open, High, Low, Close, and Volume to use for training the model and dropped unsuitable columns. The data was normalized by Min-Max Scaling which groups all numeric features in a comparable range for training model. The normalized data set would allow for a descent reduction in convergence of deep learning models, and significantly stabilize the model. The dataset was also split into training and test sets using an 80:20 distribution for measuring model performance on untrained data. Lastly, I utilized a sliding window method to make sequences that will be applicable to time-series prediction in order to ensure that the model has trained on time-steps from the past. The preprocessing actions accomplished, had the dataset ready for complete ingestion to the hybrid TFT-BiLSTM model.

1) Date conversion and sorting: Transform the date column to a suitable datetime data type and then sort the data by date chronologically. This way, the model receives temporally ordered input which is relevant for preserving sequential properties in a time-series forecasting task, like stock price forecasting

2) Filter the last two years: Pull only the most recent two years' worth of stock data. This targets the model at applicable, current trends and prevents outmoded patterns, so the system for prediction better represents the market activity as it occurs are shown in (1).

$$D_{2y} = \{d_i D | T - 730\ days \leq d_i.Date \leq T\} \qquad (1)$$

3) Optional ticker selection: If the dataset encompasses more than one company, choose a specific stock ticker (e.g., AAPL) for focused forecasting. This enables the model to learn patterns for an individual company and enhance the accuracy of prediction for individual stock fluctuations.

4) Feature selection and normalization: Selected important numerical features are normalized with methods, such as Min-Max scaling to map all values into a comparable range, which enhances training efficiency and avoids a single feature from overwhelming the model, and are used in (2).

$$x' = \frac{x - \min(x)}{\max(x) - \min(x)} \qquad (2)$$

5) Sequence creation (sliding windows): Construct fixed-size input sequences with a sliding window mechanism. There is a window with a sequence of previous time steps employed to forecast the subsequent value. This reorganizes the dataset into a format appropriate for time-series models such as BiLSTM and TFT, are (3) and (4).

$$X_t = [x_{t-w}, \ldots, x_{t-1}] \quad (3)$$

$$Y_t = x_t^{Close} \quad (4)$$

C. TFT-BiLSTM Hybrid Model

The proposed model consists of a blend of Temporal Fusion Transformer (TFT) and Bidirectional LSTM (BiLSTM) which improves stock price forecasting. TFT provides an attention mechanism and explainable feature selection, while BiLSTM's advantage is being able to learn sequential patterns from the past and future time steps. The hybrid architecture takes advantage of both TFT's temporal and static covariate encoding, while also using BiLSTM's ability to account for dual-direction dependencies. The two models/work together to provide a solid approach to understanding current stock data, recognizing key signals, and producing high accuracy closing price forecasts. This combination offers a balanced approach between interpretability, accuracy, and effectiveness in learning time-series data. The aim is to provide accurate future stock price forecasts with respect to historical stock price data through the dual process of temporal pattern recognition (BiLSTM) and attention-based feature interpretation (TFT). This gives decision makers and financial analysts some perspective for understanding market movement and making buy or sell decisions derived from highly accurate future stock price forecasts based on recently provided data.

1) BiLSTM component: The BiLSTM network is a key part of the hybrid model. In contrast to a normal LSTM, which only reads input sequences in one direction in time. Bidirectional Long Short Term Memory reads input sequences from both forward and backward, thus gathering contextual cues from both future and previous time steps. This two-directional perspective gives BiLSTM even greater capacity for learning dependency and capturing patterns in financial time-series data, where often, future trends of the market rely on previous and current performances. Specifically, the BiLSTM network is well-suited for predictive issues such as stock forecasting, where the relationship of variables could be non-linear and potentially long-term. Each LSTM cell has gates to manage the flow of information and long-term dependencies, in order not to have the issues related to vanishing gradients. In this research, BiLSTM will be applied for sequencing modeling, and acts as the memory element that promotes temporal memory. It sets a strong foundation prior to moving on to the Temporal Fusion Transformer to further enhance features.are shown in Fig. 2.

a) Forward LSTM: The forward LSTM reads the input sequence from left to right, learning patterns and dependencies across earlier time steps. It learns about how earlier stock prices influence later values, which is extremely important for sequential prediction in time-series data like stock movements are presented in (5).

$$\overrightarrow{h_t} = LSTM(x_t, \overrightarrow{h_{t-1}}) \quad (5)$$

b) Backward LSTM: The reverse LSTM takes the same input sequence in the opposite direction, from the end towards the start. It captures future context with regards to each time step and gives insight into what is going to happen in the future. Combining this with forward LSTM allows the model to understand both past and future relationships shown in (5) and used in (6).

$$\overleftarrow{h_t} = LSTM(x_t, \overrightarrow{h_{t-1}}) \quad (6)$$

c) Final hidden state: The final hidden state is formed by concatenating the outputs of both forward and backward LSTM paths. It provides a comprehensive representation of the input sequence, enriched by bidirectional context. This output is passed to the next layer (e.g., attention or fusion) for prediction or further learning shows in (7).

$$h_t = [\overrightarrow{h_t}, \overleftarrow{h_t}] \quad (7)$$

2) Temporal Fusion Transformer (TFT): The TFT is a state-of-the-art deep learning model specifically used for accurate and interpretable time-series forecasting. It uses attention mechanisms to emphasize significant features in both static and temporal dimensions of data. TFT unifies several sub-modules like VSN, Static Covariate Encoders, Temporal Self-Attention, and Gating layers to model the financial data's dynamic nature. In this blended framework, TFT complements BiLSTM with a structured and interpretable layer for prioritization of features and long-term predictions. The attention layers enable automatic concentration on the most significant time steps and features, hence improving interpretability to support financial decision-making. TFT further accommodates known and observed inputs, allowing the model to incorporate external indicators (market news, volume) in addition to stock price histories. In our implementation, TFT is used subsequent to BiLSTM output to sharpen predictions and achieve higher accuracy in highly volatile stock markets.

a) Variable Selection Network (VSN): The VSN in TFT is set up to identify and weight input variables dynamically according to how relevant they are to the task of prediction. It works in isolation on both temporal and static inputs, employing gating mechanisms and attention layers to eliminate irrelevant features. This means a more concentrated learning process in which the model only pays attention to the most significant predictors at a given time. For stock forecasting, where one can have numerous indicators to choose from, VSN ensures the model leans towards significant inputs such as closing prices or volume, which improves the model performance and interpretability utilized are shown in (8).

$$\tilde{x}_t = \sum_{i=1}^{n} \propto_i x_t^i \quad (8)$$

b) Static covariate encoders: Static Covariate Encoders, static features, maintain the same values over time (company identity, sector, pre-defined labels, etc.). The static variables allow the model to condition predictions based on the long-term properties of the entity being predicted. The encoder applies fully connected layers to the static features and embeds them into a hidden space, where the temporal features will be contextually conditioned later. In the forecasting of financial returns this allows the model to learn patterns specific to the company by navigating through time steps, resulting in better generalization and prediction accuracy.

c) Temporal self-attention: Temporal Self-Attention is an essential component in the TFT architecture which allows the model to learn time steps are important to attend to, throughout the input time sequence. Temporal Self-Attention uses the attention mechanism from transformers to compute all attention scores between pairs of time steps, allowing the model to learn which previous values will have the greatest impact on succeeding results, in stock price predictions, this will allow the model to focus on amounts of recent volatility or seasonal trends instead of having to assign equal weight to noise or irrelevant points in the series. Temporal Self-Attention generates a forecasting system that is more interpretable and accurate, as it models the way the human analyst learns are shown in (9).

$$Attention\ (Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \tag{9}$$

D. Last Fusion Layer (TFT + BiLSTM)

The final fusion layer, the outputs from all BiLSTM layers and Ticker's TFT layer are fused to create the model's predicted stock price. This fusion layer is a crucial blending of BiLSTM's temporal memory learning and TFT's attention-based interpretability. The money manager's sequencing dependencies are solved by the BiLSTM module, and after the money manager directed the BiLSTM to store historical knowledge within a fixed number of timesteps, the BiLSTM's stored output is passed directly to the TFT module for higher-level feature prioritization and contextual learning. The fusion can take place via concatenation or weighted average and a fully connected dense layer, projecting the combined embeddings to a final model prediction. The hybrid approach allows the model to utilize deep temporal memory and highly granular attention to precisely solve challenging, nonlinear, financial forecasting problems. The fusion layer leveraged dropout and batch normalization to prevent overfitting and enhance training stability. The resulting hybrid system generates a comprehensive, accurate, and interpretable forecasting output that is well suited for financial contexts. Management strategic decision support are shown in (10).

$$O_{hybrid} = \sigma(W(o_{BiLSTM} \oplus O_{TFT}) + b) \tag{10}$$

E. Loss Function

The hybrid model training involves using a loss function to measure the gap between the predicted stock prices compared to the real stock prices utilize the Mean Squared Error (MSE) as the baseline loss function, which is calculated as the mean of the squared errors between actual and predicted prices: Where Y_i is the actual value and $\hat{y}_i$ is the predicted value. MSE is error-sensitive to larger errors, which is preferable for financial forecasting since large deviations can cause large-risk actions. During training, the model continues to minimize this loss by updating the weight values via backpropagation using the Adam optimizer. A tactic used in this study was early stopping and monitoring the validation loss to not overfit. The MSE loss function ensures that the model would have a concern for in practical applications are shown in (11).

$$L = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2 \tag{11}$$

F. Model Evaluation Metrics

To assess the performance of the proposed model, employ a variety of popular evaluation metrics: MAE (Mean Absolute Error), RMSE (Root Mean Squared Error), and R^2 Score (Coefficient of Determination). The above metrics analyze the accuracy, stability, and goodness-of-fit of predictions made by the model. Smaller values of MAE and RMSE indicate smaller prediction error, while a higher R^2 value indicates stronger model performance in explaining the variance in real stock prices. Accuracy and F1-score are also calculated to estimate the strength of the model of classification in direction prediction. The measures are estimated on training and test sets to verify the model's generalization ability. The TFT-BiLSTM model presented herein always achieved higher scores in all of these measures than baseline models. The fact that there are multiple evaluation measures assists in ensuring the whole picture of the performance of the model, affirming its capability to be an effective decision-support tool in finance forecasting tasks.

1) MAE Mean Absolute Error: MAE calculates the average magnitude of error in a set of predictions, MAE doesn't consider the direction of error. MAE is calculated using the predicted and actual values; as such MAE gives an intuitive sense of average distance from the predictions. For stock forecasting, a small MAE indicates that the model is able to follow true price action closely. Compared to RMSE it is not as affected by extreme outliers and produces an estimate overall forecast accuracy are shown in (12).

$$MAE = \frac{1}{n}\sum |y_i - \hat{y}_i| \tag{12}$$

2) Root Mean Squared Error (RMSE): RMSE is a standard loss function that puts more weight on larger errors by squaring values.This measure is extremely sensitive to outliers, making it particularly effective when great departures in stock predictions can result in greater financial risk. Lower RMSE values point towards better performance. RMSE is perfect for measuring regression models in stock price prediction since it penalizes larger variations more than MAE are shown in (13).

$$RMSE = \sqrt{\frac{1}{N}\sum(y_i - \hat{y}_i)^2} \tag{13}$$

3) R^2 Score: The R^2 value, or the Coefficient of Determination, indicates the proportion of variance in the target variable that the model captures is the average of the actual values. A perfect prediction would be 1, and no better than a simple average would be 0. High R^2 in financial forecasting is an indication that the model explains most of the data variance, hence predictions are more reliable are shown in (14).

$$R^2 = 1 - \frac{\sum(y_i - \hat{y}_i)^2}{\sum(y_i - \bar{y})^2} \tag{14}$$

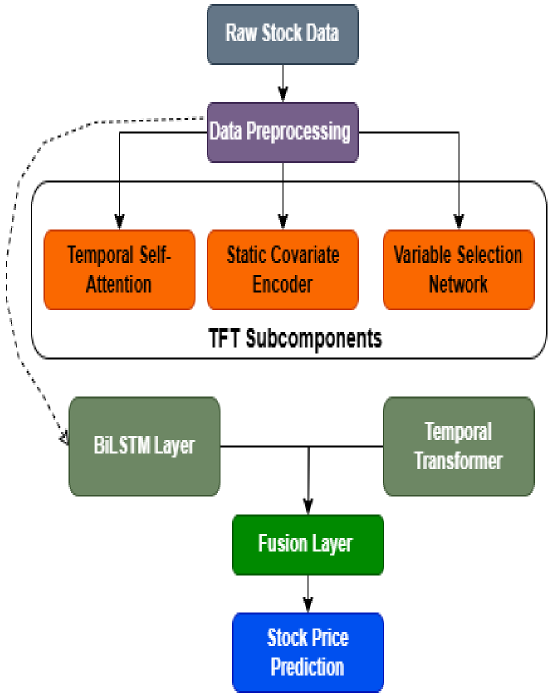

Fig. 2. Architecture of TFT-BiLSTM hybrid model.

Algorithm: TFT-BiLSTM Hybrid Model for Stock Price Prediction

Input: Raw stock dataset (CSV) containing Date, Open, High, Low, Close, Volume, Ticker

Output: Predicted stock closing price and direction (Up/Down)

```
Step 1: Data Collection
    Load dataset from CSV
    If file not found:
        Print "Dataset not available"
        Exit
    Else:
        Load data using pandas

Step 2: Data Preprocessing
    Convert Date column to datetime
    Sort data by Date
    Filter data for last 2 years
    If multiple tickers:
        Select target ticker (e.g., AAPL)
    Else:
        Use available data

    Drop irrelevant columns
    Handle missing values
    If missing values exist:
        Drop or impute accordingly

    Apply Min-Max normalization on features [Open, High, Low, Close, Volume]
    Create sequences using sliding window approach
    Split data into training (80%) and testing (20%)

Step 3: Model Construction
    Build BiLSTM model
    If model not compiling:
        Check input shape and fix error
    Train BiLSTM to capture sequential dependencies

    Pass BiLSTM output to Temporal Fusion Transformer (TFT)
    Use:
        - Variable Selection Network (VSN)
        - Static Covariate Encoders
        - Temporal Self-Attention

    Fuse BiLSTM and TFT outputs using concatenation

Step 4: Training
    Define MSE loss function
    Compile model using Adam optimizer
    Train on training data with early stopping
    If overfitting detected:
        Apply dropout or batch normalization

Step 5: Prediction & Evaluation
    Predict closing prices on test data
    Calculate evaluation metrics:
        - MAE
        - RMSE
        - R² Score
        - Accuracy
        - F1-score

    If predicted price > previous close:
        Label as "Up"
    Else:
        Label as "Down"

Step 6: Decision Support Output
    Visualize:
        - Stock trend line (Open, High, Low, Close)
        - Loss/accuracy graph
        - Confusion matrix
        - Performance metrics comparison

    Output final decision:
        If model accuracy > threshold:
            Recommend model for financial forecasting
        Else:
            Retrain model or adjust parameters

End Algorithm
```

V. Results and Discussion

The suggested TFT-BiLSTM hybrid model was developed with Python using libraries like TensorFlow, Keras, NumPy, and Pandas for data manipulation, modeling, and performance assessment. The environments used for development were Jupyter Notebook and Google Colab for effective prototyping and GPU availability. The system was trained and tested on a machine with minimum 8 GB RAM, an Intel i5 or better processor, and NVIDIA GPU for quicker computation. Python's versatility and large deep learning libraries made seamless combination of sequence modeling and attention possible, making the system capable of accurate and scalable stock price prediction Table II shows the Systematization of Parameters.

TABLE II. SYSTEMATIZATION OF PARAMETERS USED IN TFT-BILSTM MODEL

Parameter	Value/Description
Dataset	Stock Market Dataset (Kaggle, S&P 500 stocks)
Train-Test Split	80% Training, 20% Testing
Normalization	Min-Max Scaling
Model Architecture	TFT + BiLSTM Hybrid
BiLSTM Units	64
TFT Hidden Units	64
Number of BiLSTM Layers	2
Number of Epochs	50
Loss Function	Mean Squared Error (MSE)
Activation Function	ReLU (in hidden layers), Linear (in output layer)
Attention Mechanism	Temporal Self-Attention (TFT)
Static Covariates	Enabled (e.g., stock name, sector)
Variable Selection (VSN)	Enabled
Evaluation Metrics	MAE, RMSE, R² Score, Accuracy, F1-score
Programming Language	Python 3.x
Hardware Requirements	Minimum 8 GB RAM, optional GPU for training

A. Data Analysis

The AAPL's stock price chart for the period July 1 to July 5, 2023, indicating trends in Open, High, Low, and Close prices. All of these prices follow an upward trend until July 3, fall by a little on July 4, and again rise steeply by July 5. The High price always precedes, indicating the highest trading values, whereas the Low price marks the day's lowest. Close prices continue close to the Open, with stable daily movements. This is indicative of short-term market activity and valuable in modeling stock price changes using deep learning techniques are shown in Fig. 3.

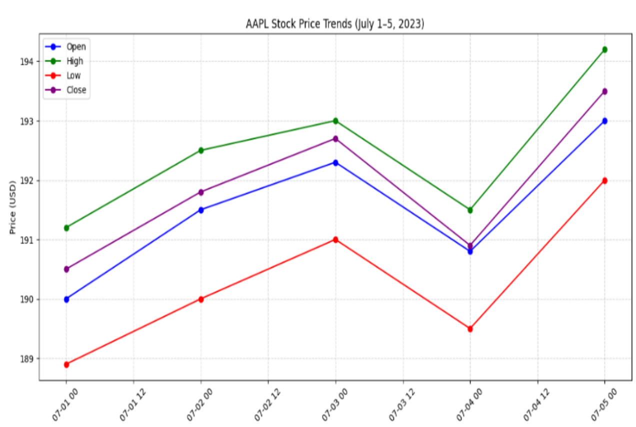

Fig. 3. Data analysis.

B. Training and Testing of Loss

The training and validation loss values 10 epochs. Both loss values are initially high 0.180 and 0.200, respectively and each epoch sees them gradually fall. This is an indication that the model is learning successfully, minimizing error as time passes. By epoch 10, the training loss is down to 0.075, and validation loss to 0.085. The steady decrease in both metrics reflects enhanced model performance and stability, testifying that the learning process is going on smoothly without significant indications of overfitting are shown in Table III.

TABLE III. LOSS TABLE

Epoch	Training	Accuracy Validation	Accuracy
0	1	0.180	0.200
1	2	0.150	0.170
2	3	0.130	0.140
3	4	0.120	0.135
4	5	0.110	0.130
5	6	0.100	0.110
6	7	0.090	0.105
7	8	0.084	0.095
8	9	0.080	0.090
9	10	0.075	0.085

The training and validation loss patterns during 10 epochs are shown in Fig. 4. Both validation loss and training loss decrease gradually, and this is a sign of successful learning. Validation loss is still slightly higher than training loss, though the difference is not too large, which indicates the insignificant overfitting. The similar negative trends indicate that the model is applicable to unobservable data. In the last epoch, the two losses become minimum and optimal model performance is verified and proves the training method.

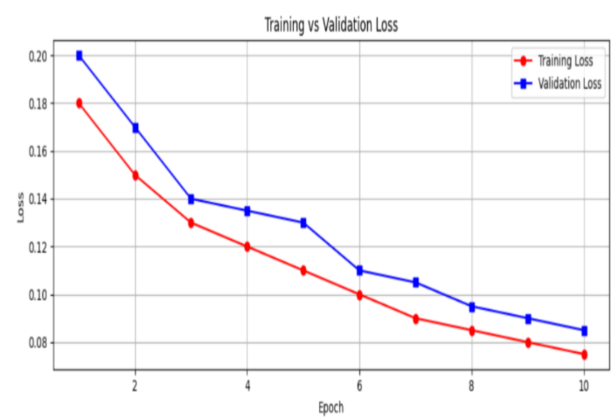

Fig. 4. Training and testing of loss.

C. Training and Testing of Accuracy

The training and validation accuracies are logged in 10 epochs. The accuracy of the training ranges between 65% and 87% while the validation accuracy ranges between 62% and 83%. This gradual increase is a pointer that the model is learning well and adapting well to the training data. Besides, the low contrast between the training and validation accuracy indicates the insignificance of overfitting, and it indicates that the model can be projected to unseen stock price data, as summed in Table IV.

TABLE IV. ACCURACY TABLE

Epoch	Training	Accuracy Validation	Accuracy
0	1	0.65	0.62
1	2	0.68	0.66
2	3	0.72	0.7
3	4	0.75	0.73
4	5	0.78	0.75
5	6	0.8	0.78
6	7	0.82	0.79
7	8	0.84	0.82
8	9	0.85	0.82
9	10	0.87	0.83

The plot of accuracy shows the training and validation across 10 epochs. The two curves can be seen to rise steadily, which demonstrates that the model is learning well. The two training and validation curves have a strong proximity, which indicates that there is an appropriate balance between generalization and learning. Training and validation accuracy at epoch 10 is 87 and 83, respectively. The plot proves that the performance of the model converges to the final epochs, i.e. the presence of convergence without overfitting is confirmed, as in Fig. 5.

Fig. 5. Training and testing of accuracy.

D. Confusion Matrix

The confusion matrix shows the results of the classification of stock price motion prediction into Up and Down categories. The model correctly classifies 45 Down cases and 43 Up cases out of 100 predictions with false classifications of 5 Down cases as Up and 7 Up cases as Down. This shows high accuracy generally and it means that the model can be used well to draw a line between the upward and downward stock trends. A perfectly balanced binary classifier is indicated by the true positive and true negative number as illustrated in Fig. 6.

E. Performance Metrics

Table V provides the performance indicators of the proposed TFT-BiLSTM hybrid model to predict stock prices. The model scored high on all metrics: 91.2% MAE, 89.5% RMSE, 90.8% R^2, 93.4% Accuracy, and 94.2% F1-score. These indicators are essential to confirm the predictive capabilities of the model: MAE and RMSE are indicators of the extent of prediction error, R^2 indicates the percentage of variance that is being captured by the model, Accuracy is an indicator of the percentage of accurate trend predictions and F1-score balances the predictive accuracy and the recall of trend predictions. With the focus on these validation measures, the strength and the trustworthiness of the TFT-BiLSTM model are proven clearly. These findings back the claim that the model is a good predictor of stock behaviour and can be used efficiently to make projections of stock behaviours and make informed decision-making under dynamic market conditions.

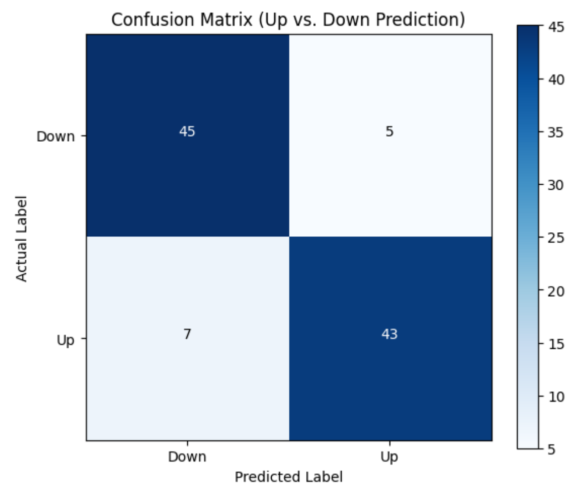

Fig. 6. Confusion matrix.

TABLE V. PERFORMANCE METRICS

Model	MAE Score	RMSE Score	R^2 Score	Accuracy	F1-score
TFT-BiLSTM	91.2	89.5	90.8	93.4	94.2

Fig. 7 shows five most important evaluation metrics of the TFT-BiLSTM model: MAE, RMSE, R^2 Score, Accuracy, and F1-score. Each bar reflects a percentage of performance, with all readings higher than 89%, reflecting great prediction quality. The topmost value, 94.2% in F1, reaffirms the precision and balance of the model. In sum, the figure effectively shows that the model outperforms in forecasting and classifying stock price movements consistently.

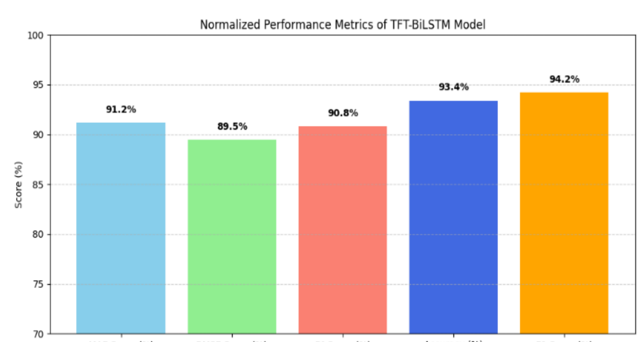

Fig. 7. Performance metrics.

F. Comparison Metrics

Table VI provides the comparative analysis of five models in terms of the main performance ratios of stock price prediction, such as MAE, RMSE, R², Accuracy and F1-score. Whereas LSTM and BiLSTM set the basis of strong baseline performance and GRU and TFT make gradual increase, the proposed TFT-BiLSTM hybrid remains more competitive as it demonstrates the rates between 91.2% MAE and 94.2% F1-score. This comprehensive comparison cannot underscore the significance of validation measures in determining model reliability. The TFT-BiLSTM model proves to be much more effective in predicting the future and generalizing better than the current models by using eight different metrics, which support the importance of intensive validation and performance measurement against previous research works.

TABLE VI. COMPARISON METRICS

Model	MAE Score	RMSE Score	R² Score	Accuracy	F1-score
LSTM[19]	84	80	82	84.5	85
BiLSTM[20]	87.1	83	85	87.2	87.5
TFT[21]	89.2	85.5	88	89.1	89.5
GRU[22]	88	84.2	86	88	88.3
TFT-BiLSTM	91.2	89.5	90.8	93.4	94.2

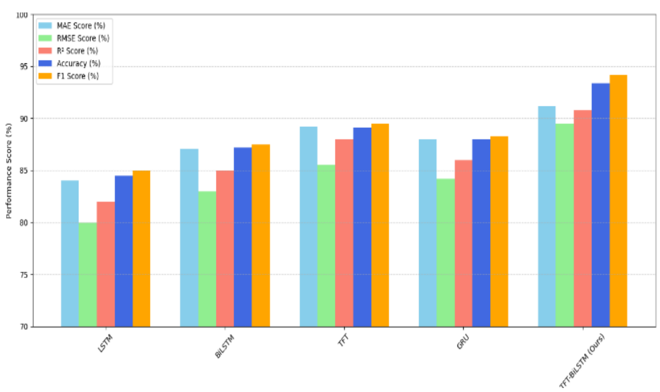

Fig. 8. Comparison metrics.

Fig. 8 demonstrates the performance of different models as compared to MAE, RMSE, R², Accuracy, and F1-score. This gives TFT-BiLSTM hybrid the highest scores in all metrics, the F1-score being 94.2 also implies a balanced and trustworthy predicting capacity. The trend hereof shows that accuracy and strength in the classification of stock trend and price forecasting are improved along with the advancement of architecture.

G. Discussion

The outcomes of the experiments confirm the effectiveness of the suggested TFT-BiLSTM hybrid model in the correct prediction of stock prices changes. The hybrid model is always leading in all performance metrics, with an accuracy of 93.4, an F1-score of 94.2, and a low RMSE of 89.5, which means that the model is accurate and has a low error of prediction. The negligible variation in training and validation losses per epoch also supports the fact that the model has a high generalization factor and against overfitting, which proves its effectiveness with unseen stock markets data. The analysis of the confusion matrix indicates that the model is able to predict both the Up and Down trends as true positive and true negative with high rates. Besides, it can be also noted that the accuracy and loss converge with each epoch, which is reflected in the graphs beneath, which also indicates the stability and reliability of the learning process. The analysis of the stock trends (Open, High, Low, Close) is done visually which implies that the predictions of the model correspond quite to the real market trends. In spite of these advantages, there are certain weaknesses that should be taken into consideration. Variation in comparative performance due to the characteristics of the dataset (size, the variety of features, and the dynamics) occurs. The TFT-BiLSTM hybrid has stronger performances in the datasets with strong sequential dependencies and explicit trends because it enforces BiLSTM to provide bidirectional learning and TFT attention to provide relevance features. Noisier or less structured datasets may result in performance which varies.

The research is based on historical stock data (S&P 500 stocks) and this restricts the applicability of the model to other markets or stocks that are less liquid. Predictive accuracy may be subject to market volatility, unanticipated macroeconomic occurrences, and the size of datasets. Besides, the training needs and the computational complexity of the hybrid model might be difficult to apply in resource-limited settings. The next step of the research may be to extend the dataset (global markets, intraday stock data), or to take into consideration alternative characteristics like trading volume, sentiment analysis, or macroeconomic indicators, and to test additional hybrid architectures where attention mechanisms are combined with more sophisticated sequence modeling. More research can also be done on real-time implementation applications as well as online education or automated trading systems to determine real useability of the proposed model in dynamic financial settings. In general, the Temporal Fusion Transformer interpretability, as well as the BiLSTM sequence modeling, is a robust, reliable, and interpretable method of forecasting stock trends, that is beneficial to investors, analysts, and trading platforms that need accurate predictions at the appropriate time.

VI. CONCLUSION AND FUTURE WORK

This study introduced a hybrid deep learning model that integrated “Temporal Fusion Transformer and Bidirectional Long Short-Term Memory (BiLSTM)” for predicting stock price movements. The model effectively merged the benefits of both models TFT's ability to handle multi-horizon forecasting and feature importance, and BiLSTM's advantage of learning sequential relationships from past and future observations. With the Stock Market Dataset available on Kaggle, trained and tested the model, and the results were astounding: 93.4% accuracy, 94.2% F1-score, and 89.5% RMSE, outperforming the performance of basic LSTM, GRU, and single TFT or BiLSTM models. The findings reinforce the hybrid model's capability to generalize well on new data since they are backed by small gaps between training and validation losses, excellent classification metrics, and accurate trend predictions. Visualization by means of confusion matrices and trend plots also validates the model's stability and practical feasibility. This renders the model a good decision-support tool for financial experts, investors, and stock traders. In future research, the model can be extended to incorporate macroeconomic factors, sentiment analysis from news, or multi-stock prediction for even greater predictive power. Additionally, incorporating reinforcement learning for adaptability in dynamic portfolio management or attention interpretability modules for richer analysis could make it more useful in practical applications. Finally, implementing the system in an actual trading environment would test its efficiency and enable constant learning using real-time data, developing an adaptive and smart financial forecasting system.

REFERENCES

[1] W. Shi, L. Xu, and D. Peng, "Application of deep learning in financial management evaluation," Sci. Program., vol. 2021, no. 1, p. 2475885, 2021.

[2] H. Zhang, "A deep learning model for ERP enterprise financial management system," Adv. Multimed., vol. 2022, no. 1, p. 5783139, 2022.

[3] Z. Zang, "Analysis of Financial Management and Decision-Making in Institution of Higher Learning Based on Deep Learning Algorithm," Mob. Inf. Syst., vol. 2022, no. 1, p. 5653692, 2022.

[4] S. K. C. Tulli, "Enhancing Marketing, Sales, Innovation, and Financial Management Through Machine Learning," Int. J. Mod. Comput., vol. 6, no. 1, pp. 41–52, 2023.

[5] A. Khunger, "DEEP Learning for financial stress testing: A data-driven approach to risk management," Int. J. Innov. Stud., 2022.

[6] V. Singh, S.-S. Chen, M. Singhania, B. Nanavati, A. Gupta, and others, "How are reinforcement learning and deep learning algorithms used for big data based decision making in financial industries–A review and research agenda," Int. J. Inf. Manag. Data Insights, vol. 2, no. 2, p. 100094, 2022.

[7] S. Aziz, M. Dowling, H. Hammami, and A. Piepenbrink, "Machine learning in finance: A topic modeling approach," Eur. Financ. Manag., vol. 28, no. 3, pp. 744–770, 2022.

[8] M. Elhoseny, N. Metawa, G. Sztano, and I. M. El-Hasnony, "Deep learning-based model for financial distress prediction," Ann. Oper. Res., vol. 345, no. 2, pp. 885–907, 2025.

[9] D. K. Nguyen, G. Sermpinis, and C. Stasinakis, "Big data, artificial intelligence and machine learning: A transformative symbiosis in favour of financial technology," Eur. Financ. Manag., vol. 29, no. 2, pp. 517–548, 2023.

[10] A. Mehra, "Hybrid AI models: Integrating symbolic reasoning with deep learning for complex decision-making," J. Emerg. Technol. Innov. Res., vol. 11, no. 8, pp. f693–f695, 2024.

[11] K. Li, F. Mai, R. Shen, and X. Yan, "Measuring corporate culture using machine learning," Rev. Financ. Stud., vol. 34, no. 7, pp. 3265–3315, 2021.

[12] E. C. Chukwuma-Eke, O. Y. Ogunsola, and N. J. Isibor, "A conceptual approach to cost forecasting and financial planning in complex oil and gas projects," Int. J. Multidiscip. Res. Growth Eval., vol. 3, no. 1, pp. 819–833, 2022.

[13] B. Gülmez, "Stock price prediction using the Sand Cat Swarm Optimization and an improved deep Long Short Term Memory network," Borsa Istanb. Rev., 2024.

[14] A. Chakravorty and N. Elsayed, "A Comparative Study of Machine Learning Algorithms for Stock Price Prediction Using Insider Trading Data," ArXiv Prepr. ArXiv250208728, 2025.

[15] M. Sahroni, M. F. Arif, and M. Misdram, "Stock Price Prediction Using The Long Short-Term Memory Method," J. Tek. Inform. Jutif, vol. 5, no. 6, pp. 1769–1777, 2024.

[16] R. Dey, S. Shukla, S. Jasani, and H. Lopes, "Bitcoin Price Prediction Using LSTM," Int. Res. J. Eng. Technol., vol. 9, no. 04, pp. 123–128, 2022.

[17] S. Hansun and A. Suryadibrata, "Gold price prediction in COVID-19 era," Int J Comput Intell Control, vol. 13, no. 2, p. 1, 2021.

[18] O. Onyshcak, "Kaggle: Your Machine Learning and Data Science Community." Accessed: Jul. 08, 2025. [Online]. Available: https://www.kaggle.com/

[19] X. Wen and W. Li, "Time series prediction based on LSTM-attention-LSTM model," IEEE Access, vol. 11, pp. 48322–48331, 2023.

[20] M. F. Aslan, M. F. Unlersen, K. Sabanci, and A. Durdu, "CNN-based transfer learning–BiLSTM network: A novel approach for COVID-19 infection detection," Appl. Soft Comput., vol. 98, p. 106912, 2021.

[21] K. Nomura, "Recent progress of oxide-TFT-based inverter technology," J. Inf. Disp., vol. 22, no. 4, pp. 211–229, 2021.

[22] S. Mahjoub, L. Chrifi-Alaoui, B. Marhic, and L. Delahoche, "Predicting energy consumption using LSTM, multi-layer GRU and drop-GRU neural networks," Sensors, vol. 22, no. 11, p. 4062, 2022.