

A Bidirectional LSTM–Sentiment Fusion Framework for Dynamic Financial Market Prediction

Minal Dhankar, Neha Gupta
School of Computer Applications,
Manav Rachna International Institute of Research & Studies,
Faridabad, India

Abstract—Financial market prediction can be said to be a great challenge because of the intrinsic fluctuation, non-stationarity and multi-faceted influence of the economic indicators, world events, as well as the voter sentiment. Conventional models can easily miss the time dependence and emotional aspects inherent in market data, and the results in poor forecasting precision. The paper presents a sequence-based modelling with sentiment analysis based on textual information like news articles and social media, which incorporates a two-way LSTM-sentiment fusion framework. It discloses that sentiment integration discerns as well as polishes predictive results into alignment with temporal characteristics with real-time emotive drivers.

Keywords—*Bidirectional LSTM; deep learning; financial market; modelling; sentiment analysis*

I. INTRODUCTION

The stability of market trends is one of the pillars of stability in the economy, investment strategy, and the policies followed in the world of finance, which is ever dynamic and evolving. The financial markets, with their complex demand and supply relationships, geopolitical context, and macroeconomic measures, demonstrate both astonishing and frustrating trends [1]. The role of the correct prediction should not be overestimated; it will enable investors to mitigate the effect of risks, as well as corporations to organise the capital spending more efficiently, and enable governments to avoid crises. The discipline has since received indefatigable scholarly attention and has grown over the decades to encompass both more primitive statistical models as well as more complicated computational paradigms [2]. Time-series models, e.g., autoregressive integrated moving average models, were initially capable of modelling a linear relationship between the past data, but did not work when non-linear volatilities arose. Machine learning could introduce a new period, and approaches such as support vector machines and random forests demonstrated more successful outcomes in the work with complex data. More recently, deep learning networks, like recurrent neural networks, are popular because of their capability to handle sequential data, and therefore they can make subtle predictions of the price of

stocks, changes in the exchange rate, and the movement of commodities [3].

With such actions, the old-fashioned models of forecasting usually fail to deal with the complexity of market forces, where figures can be used to paint a partial picture. Recurrent neural networks are highly successful in identifying patterns in time; however, they also tend to fail to capture the psychological forces in terms of market sentiment, i.e. investor optimism or panic in response to specific events occurring in the market. The outcome of such negligence is robust exercising in stable situational conditions and feeble in turbulent situational conditions, like recessions or political crises [4]. Moreover, standard recurrent models are only able to process data in a single direction, which limits their ability to contextualise future values by examining past sequences, which can contain bidirectional dependencies, which are inherent to financial time series [5]. Questions on how to incorporate the elements of qualitative nature exist: how can we quantify and combine emotional colours present in textual sources without noise? However, there are methods of hybrid models, in which sentiment analysis is attempted to be coupled to neural networks, though these are largely quite superficial in the methods they apply [6]. The nature of hybrid methods is that they take into account only simple polarity scores that are not the intensity of the popular discourse, but only its context specificity. It means there is a clear literature gap- it would be desirable to have frameworks which are not merely bidirectional in terms of observing the past prices but which are also sensitive to the sentiment dynamics in order to enhance predictive fidelity. The disjunction will prompt the community to stop the silo-based analysis of quantitative data and act as a reminder to the community that interdisciplinary fusion has been the enabling factor in the past of financial econometrics [7].

II. LITERATURE REVIEW

The subsequent works (summarised in Table I) outline the developments of LSTM-based models, attention, reinforcement learning, embeddings, as well as multi-source fusion without suggesting the obstacles like data imbalance, market volatility, and model generalizability.

TABLE I. RELATED WORKS

Author–Year	Paper Title	About the Paper	Methodology Used	Limitations
Fu et al., 2019 [1]	Sentiment-Aware Trading Volume Prediction for P2P Market Using LSTM	Predicts P2P trading volume incorporating investor sentiment correlations.	TextCNN for sentiment, LSTM for volume prediction, causality analysis.	Limited to P2P domain; comments may not represent full market; older dataset reduces applicability today.
Zhang et al., 2022 [2]	Decision Fusion for Stock Market Prediction: Systematic Review	Reviews ensemble and decision-fusion models for stock prediction.	Systematic review of fusion techniques, classification of ensemble methods.	No experimental validation; fusion taxonomy may become outdated given fast model evolution.
Yekrangi & Nikolov, 2023 [3]	Domain-Specific Sentiment Analysis for Financial Markets	Evaluates embeddings and classification algorithms for financial sentiment. Proposes optimized embedding layer.	LSTM, CNN, SVM, RNN, fine-tuned embeddings (CBOW, GloVe, BERT).	Context-dependency issues persist; sentiment ambiguity; limited financial-domain labeled datasets.
Mu et al., 2023 [4]	Stock Price Prediction Using Investor Sentiment and Optimized Deep Learning	Integrates sentiment index with SSA-optimized LSTM for Chinese stock prediction.	SSA-optimized LSTM, custom sentiment dictionary, multi-source data fusion.	Sentiment dictionary manually built; model tailored for China’s market; short-term accuracy higher than long-term.
Girsang & Stanley, 2023 [5]	Hybrid LSTM-GRU for Cryptocurrency Price Forecasting Using FinBERT	Predicts ETH and SOL using combined price + social media sentiment.	FinBERT sentiment scoring, Hybrid LSTM-GRU.	Crypto markets highly volatile; overreliance on Twitter sentiment; may fail during extreme market shocks.
Peivandizadeh et al., 2024 [6]	Stock Market Prediction With Transductive LSTM and Social Media Sentiment Analysis	Combines social media sentiment with market data to improve stock prediction. Addresses data imbalance in sentiment classes.	Transductive LSTM, Off-policy PPO for imbalance handling, ablation studies.	Sensitive to class imbalance; social-media noise may affect accuracy; requires extensive tuning for reinforcement learning.
Bacco et al., 2024 [7]	Stock Prediction Using LSTM and Tweet Sentiment Under High Uncertainty	Studies bank stock prediction amid geopolitical and macroeconomic shocks using financial + sentiment data.	LSTM, FinBERT sentiment extracted from Twitter, multi-market comparison (US/EU).	High volatility reduces reliability; market differences limit generalizability; sentiment may not represent institutional behavior.
Alam et al., 2024 [8]	Enhancing Stock Market Prediction: LSTM-DNN on 26 Datasets	Builds a robust hybrid LSTM-DNN model validated on 26 stock datasets.	Hybrid LSTM-DNN, ablation studies, multi-company evaluation.	Uses only structured price data (no sentiment); may overfit despite robustness; lacks real-time testing.
Imanto et al., 2025 [9]	Fin-ABSA Using Fusion-Guided Attention LSTM and Graph Convolutional Networks	Improves aspect-based sentiment analysis for financial texts by combining syntactic and semantic cues.	Attention LSTM + GCN fusion model (FGAG).	Requires high-quality dependency parsing; computationally expensive; domain-specific aspects may not transfer across markets.
Babak Amiri, Amirali Haddadi & Kosar Farajpour Mojdehi (2025) [10]	A Novel Hybrid GCN-LSTM Algorithm for Energy Stock Price Prediction: Leveraging Temporal Dynamics and Inter-Stock Relationships	Proposes a hybrid deep-learning model combining GCN and LSTM with attention to improve prediction of highly volatile energy stock prices by modeling both temporal trends and inter-stock relationships.	Dynamic Time Warping (DTW) to construct correlation matrix. Graph Convolutional Network (GCN) to extract spatial dependencies. Attention-enhanced LSTM to capture temporal patterns. Comparative evaluation versus GRU, MLP, LR, LSTM	Performance drops in highly volatile or sudden-shock conditions. Model does not include external sentiment/news data. Limited testing outside the energy sector. Depends heavily on historical price data only
Nejatbakhsh & Aliasgari, 2026 [11]	Enhancing Stock Market Prediction With Hybrid Deep Learning: LSTM, Transformer Attention, Federated Learning, and Sentiment Analysis	Develops a hybrid architecture integrating deep learning and privacy-preserving FL for stock forecasting.	LSTM + Transformer attention, FinBERT sentiment, Federated Learning simulation.	Computationally heavy; federated setup may limit real-world deployment; requires consistent data quality across nodes.

III. THEORETICAL BACKGROUND

The operational basis of prediction in financial markets is based on two fundamental areas, including mathematical properties of financial time series and the computational design principles of deep-learning-based sequence models. These areas can be learned, and then it is possible to examine the performance of a bidirectional LSTM-sentiment fusion model within dynamic markets. Financial systems cannot be associated with relatively simple numbers; they contain structural anomalies such as non-stationarity, fat-tailed distributions, volatility clustering, behavioural bias, and nonlinear dependence. Recurrent deep learning models have mechanisms

for extracting temporal structure that are not exploited by conventional statistical models [8].

A. Financial Time Series Theory

Another well-known fact about financial time series is that they are non-stationary, i.e. their statistical behaviour, i.e. mean, variance and autocorrelation, changes with time. Non-stationarity is likely to be induced by macroeconomic events, structural breaks and regime shifts [9]. As a result, the models, which assume the existence of stable distributions, are poor performers. Neural networks (in particular, LSTMs) deal with these issues since they learn representations dynamically without taking a strict stationarity condition [10].

The other important behaviour is the volatility clustering, whereby the volatility cases tend to be followed by high volatility, and low volatility tends to be followed by low volatility. The occurrence of such a phenomenon implies that there is market uncertainty, and it plays a vital role in the day-to-day forecasting. Young sequences are somehow advantageous in deep-learning models as they locate time trends within sequences of such clusters.

Financial information is no exception: it also has heavy-tailed distributions, i.e. extreme price changes are more widespread than should have been the case assuming it was distributed normally. That is because the machine-learning architectures can handle the heavy-tailed behaviour as compared to the statistical models, since they do not have distributional constraints [11].

B. Deep Learning Background

Long Short-Term Memory (LSTM) networks were invented in an attempt to address such inadequacies. The LSTMs have been created based on the inclusion of three primary gates, such as forget gate, input and output gates, and an internal memory cell in which long-range dependencies are stored [12].

The equations of the standard LSTM are:

Forget Gate:

$$f_t = \sigma(W_f[h_{t-1}, x_t] + b_f) \quad (1)$$

Input Gate:

$$i_t = \sigma(W_i[h_{t-1}, x_t] + b_i) \quad (2)$$

Candidate Memory:

$$C_t = \tanh(W_c[h_{t-1}, x_t] + b_c) \quad (3)$$

Cell State Update:

$$C_t = f_t \cdot C_{t-1} + i_t \cdot C_t \quad (4)$$

Output Gate:

$$o_t = \sigma(W_o[h_{t-1}, x_t] + b_o) \quad (5)$$

Hidden State:

$$h_t = o_t \cdot \tanh(C_t) \quad (6)$$

All these equations enable LSTMs with the capability to learn sequence-long term relations, thus it is most suitable for financial series that need past volatility regime and behaviour swings memory [13].

C. Bidirectional LSTM Theory

It is a mechanism that takes data forward and backwards. It is divided into two parallel layers of LSTM:

One reads the sequence from $t = 1 \rightarrow T$

The other reads from $t = T \rightarrow 1$

For each timestep:

$$h_t^{forward} = LSTM_{forward}(x_t) \quad (7)$$

$$h_t^{backward} = LSTM_{backward}(x_t) \quad (8)$$

These are concatenated:

$$H_t = [h_t^{forward}, h_t^{backward}] \quad (9)$$

Bidirectionality not only covers structural delays but also matches volatility-diffusion behaviour, and rebates the effects of sentiment spikes through contextualization by forward information [14].

IV. METHODOLOGICAL FRAMEWORK

After the data preparation, the analysis methods are focused on the BiLSTM structure with sentiment fusion. LSTM is a two-way iteration on financial time-series data, each input sample is a fixed-length window of $T = 30$ consecutive trading days, containing OHLCV features (5 price/volume variables) and the aligned sentiment score, yielding an input tensor of shape $[30 \times 6]$ per sample [15]. Vectors features (Sentiment features) are combined in the hidden layer of the LSTM via concatenation or attention to allow the model to adjust its prediction according to the current psychology in the market. The training process includes a division of the dataset into time intervals of training with and without validation in time, the use of adaptive learning rates of backpropagation to minimise the weights, and the use of regularisation to avoid overfitting. Measures of evaluation are aimed at predictive quality, which determines the forecast of the framework by minimising errors and robustness tests in the different market conditions. The flowchart below is an analytical pipeline visualising the sequence of steps involved in processing raw data to the final prediction output in an iterative manner.

A. Design of the Sentiment Analysis Module

Sentiment scoring in this framework uses FinBERT (ProsusAI/finbert), a BERT-based model pre-trained on financial corpora and fine-tuned on the Financial PhraseBank dataset to classify text as positive, negative, or neutral. Since the study uses synthetic data, FinBERT scoring is simulated by generating continuous polarity values in the range $[-1, +1]$ using a correlated random process ($\rho = 0.8$ with next-period returns), with Gaussian noise added ($\sigma = 0.1$) to replicate realistic scoring variance. In a real deployment, raw text from news or social media would be passed through FinBERT's classifier to produce these scores directly [16].

One of the initial processes that entails developing the sentiment module is the identification of valuable sources of sentiment information. In the real-deployment design of this framework, three primary data sources are targeted: (1) Financial news headlines from Reuters and Bloomberg financial feeds, covering the period January 2018 to December 2022 (approximately 5 years, aligned with the 5,000 synthetic trading days); (2) Social media posts from Twitter/X (via the Academic Research API), filtered by equity-related cashtags (e.g., \$SPY, \$AAPL); and (3) Analyst reports and earnings call transcripts from SEC EDGAR filings. Since this study uses synthetic data, these sources are not scraped live — instead, the FinBERT scoring behaviour is replicated numerically. In a real deployment, all three streams would be passed through FinBERT (ProsusAI/finbert) to generate per-document polarity scores.

The sentiment extraction pipeline for real data would follow these steps: (1) Raw text is collected and timestamped at the

document level. (2) Each document is tokenised using FinBERT's WordPiece tokeniser with a maximum sequence length of 512 tokens. (3) The model outputs a three-class probability distribution (positive, negative, neutral); the final sentiment score is computed as: $\text{Score} = P(\text{positive}) - P(\text{negative})$, yielding a value in $[-1, +1]$. (4) Multiple documents within the same trading day are aggregated by weighted average, with weights proportional to source credibility (news > analyst report > social media). (5) If no document exists for a given day, the previous day's score is carried forward (forward-fill), matching the handling used for missing financial data.

Both kinds of sources present distinct behavioural cues, as news headlines will present sentiment impacting highly and with formalised linguistic structures, whereas the social media posts will present the raw and immediate responses of retail investors. Since real textual data was not available for this study, synthetic sentiment scores were generated using the following procedure: (1) A base sentiment signal was constructed as a mean-reverting process centred at zero, using an AR(1) process: $S_t = 0.7 \cdot S_{t-1} + \varepsilon_t$, where $\varepsilon_t \sim N(0, 0.1)$. (2) The signal was then correlated with next-period log returns at $\rho = 0.8$ by blending with the return series. (3) Gaussian noise ($\sigma = 0.05$) was added to simulate real-world scoring variance and prevent trivial leakage. (4) The final scores were clipped to $[-1, +1]$ to match FinBERT's output range. This design ensures that sentiment has genuine predictive power (justifying the performance gains), while the added noise prevents the model from simply memorising the sentiment-return mapping [17].

Lastly, the sentiment module would be a part of the deep-learning pipeline. The sentiment sequence will be processed as a parallel time-series input, i.e. it will be normalised, windowed and put together in a batch just as is the case with the price data. Normalisation brings about consistency in the numbers so that the sentiment values are brought into an even scale. The sentiment information gets organised into overlapping sequences by windowing, and the BiLSTM identifies attention to behavioural relationships across time. Incorporated into the fusion layer is the combination of the sentiment outputs and the bidirectional price outputs, which are then merged to provide a single representation. The sentiment module is also applied to give the structural signals an emotional context at that stage, which better predicts the continuation or reversal of the trend in a given model [18].

B. Fusion Strategy

In the present case, late fusion offers a theoretical background of how behavioural and structural information may be fused together. The BiLSTM derives the forward and backward time relationships out of the price sequence. At the same time, the sequence of sentiments also receives its encoding procedure that provides a map of variations in behaviour onto a latent space. Upon the congruence of both streams at the fusion layer, the streams are combined to form a joint representation [19].

The mathematical model of the fusion approach may be presented in the following way. Let:

$$H_t^{\text{price}} \quad (10)$$

Represent forward and backwards hidden state pairs of the Bidirectional LSTM, and

$$S_t \quad (11)$$

denote the sentiment vector at time t

The fusion reaction may be defined as:

Fusion Equation:

$$F_t = \alpha H_t^{\text{price}} + (1 - \alpha) S_t \quad (12)$$

α is a parameter that is learned and which dictates the relative importance of structural and behavioural signals. The model focuses on price-based patterns when there is stability in the conditions of the market. In the time of high volatility or sentiment, $(1 - \alpha)$ Rises and heightens the part played by sentiment. Such dynamic adaptation is reflected in actual financial situations in which behavioural signals tend to become prominent in uncertain situations.

The power of the fusion mechanism is that the model can control its decision-making pathways. When the price-based signals are not in differentiation, behavioural pressure can be decoded by sentiment features. On the other hand, when sentiment is noisy, structural stability is recovered because of the bidirectional temporal context. That balancing mechanism enables the model to identify the reversal of trends earlier, identify the change of momentum, and address prediction drift [20].

Flowchart (Fig. 1) starts by merging parallel nodes in the data collection streams, which are in the form of financial and sentiment data, into preprocessing, which is then fed into the BiLSTM model containing sentiment fusion, and then training, evaluation, and prediction stages. Every node depicts a major step, and the arrows show the flow of data and where decisions can be made to be reiterated. To provide an elaborate elaboration, the flow-chart begins with the "Start node", which branches into the collecting financial data and upcoming sentiment data through the extension of the line into the node of "Collect Financial Data" and "Collect Sentiment Data" respectively, to accentuate the bilateral aspect of the input of the financial data and that of the sentiment data respectively. These streams are combined at the "Preprocess Data" node, which does some normalisation of financial input, generates sentiment, and generates refined features. That will result in the Build BiLSTM Model node, in which stacking of bidirectional layers is done to handle sequences in both directions, with respect to refining feature-extracting capabilities by having the ability to examine historical events as well as projecting events in advance. Emotional vectors are incorporated in the "Fuse Sentiment" node and, through the application of emotional presentation techniques such as weighted summation, are used to modify LSTM gates to further refine the hidden states to make more subtle predictions. The flow continues to the next train model with loss computation and optimisation loops, and a decision branch that converts the convergence checks, which could take one more loop back to another epoch. When it has been satisfactorily trained, it proceeds to Evaluate Model and compares with holdouts and ends ultimately at the generate Predictions node to make projections in a dynamic market. The form emphasises the modularity in the framework for repeated

improvements and the benefits of a bidirectional processing to reduce the shortcomings of unidirectional models through the provision of a full picture of the time series and the reduction of gaps in the analysis with only quantitative data through the fusion of sentiments.

In order to explain the methodology further, tables are included that would order the elements. Table II provides a list of possible data collection sources, splitting financial data into columns based on asset type, data source, and time-granularity. Table III describes the preprocessing methods, and the rows consist of steps to preprocess, e.g. scaling, tokenisation,

alignment, and the columns contain the input type (financial, sentiment), the method, and the purpose, such as min-max scaling to ensure values from the output of the data used to train the model are constrained between zero and one. Analysis methods are outlined in Table IV, organised into analysis stages such as model building, fusion and evaluation, divided by sub-columns into the specifics of the BiLSTM (e.g. the layer structure of forward and backward passes), methods of sentence sentiment integration (e.g. embedding concatenation), and validation techniques (e.g. cross-validation over time-series to maintain the order of sequence) [21].

TABLE II. DATA COLLECTION SOURCES AND CHARACTERISTICS

Data Type	Asset/Source Category	Primary Sources	Temporal Granularity	Coverage Period	Purpose
Financial Time-Series	Equities	Stock exchange APIs, financial databases	Daily / Intraday	Multi-year (e.g., 5–10 years)	Capture price and volume trends
Financial Time-Series	Forex	Currency exchange APIs	Hourly / Daily	Multi-year	Model exchange rate fluctuations
Financial Time-Series	Cryptocurrencies	Major exchange feeds (e.g., Binance, Coinbase)	Minute / Hourly	Multi-year	Handle high-volatility assets
Sentiment Text	Financial News	News aggregators, financial portals	Real-time / Daily	Aligned with financial data	Extract market-moving events
Sentiment Text	Social Media	Public posts via APIs/scrapers	Real-time	Aligned with financial data	Capture retail investor sentiment

TABLE III. PREPROCESSING TECHNIQUES.

Step	Input Type	Technique / Method	Parameters / Details	Purpose
Missing Value Handling	Financial	Linear interpolation, forward/backward fill	Applied when gaps < 5% of sequence	Ensure continuity in time-series
Normalization	Financial	Min-Max scaling or Standard scaling	Scale to [0,1] or zero mean, unit variance	Stabilize gradient flow in LSTM training
Stationarity Check	Financial	Differencing (first-order), ADF test	Differencing applied if non-stationary	Reduce trend and seasonality effects
Tokenization	Sentiment Text	Word-level tokenization	Using pre-trained tokenizer	Prepare text for embedding
Stop Words Removal	Sentiment Text	Removal of common English stop words	Standard library lists	Reduce noise in sentiment scoring
Sentiment Scoring	Sentiment Text	FinBERT (ProsusAI / finbert, fine-tuned on financial PhraseBank)	Polarity (-1 to +1) + intensity score	Quantify emotional tone and strength
Temporal Alignment	Both	Timestamp matching and nearest-neighbor join	Align sentiment to closest financial timestamp	Enable accurate fusion of modalities

TABLE IV. ANALYSIS TECHNIQUES AND MODEL COMPONENTS

Phase	Component / Technique	Key Details / Configuration	Role / Advantage
Sequential Modeling	Bidirectional LSTM	2 stacked BiLSTM layers, 128 hidden units per direction, dropout 0.3; input window T = 30 timesteps, input features = 6 (5 OHLCV + 1 sentiment)	Captures past and future context in sequences
Sentiment Fusion	Feature-level fusion (concatenation / attention)	Sentiment embedding concatenated to LSTM hidden states or used via attention gate	Integrates market psychology into price prediction
Model Training	Backpropagation through time (BPTT)	Adam optimizer, lr = 0.001 with patience-based decay (factor 0.5), batch size 64, max 100 epochs, early stopping (patience 10)	Efficient weight updates on sequential data
Regularization	Dropout, L2 regularization	Dropout rate 0.2–0.5, L2 penalty on weights	Prevent overfitting on noisy financial data
Validation Strategy	Time-series cross-validation / Walk-forward	Chronological splits, no random shuffling	Respect temporal order, realistic out-of-sample test
Evaluation Metrics	RMSE, MAE, Directional accuracy, Sharpe ratio (simulated)	Primary: error-based; Secondary: financial utility	Assess prediction accuracy and practical value
Scenario Testing	High-volatility periods, event-driven windows	Subset evaluation during earnings, policy announcements	Validate robustness during market stress

The exact configuration used in this study is as follows: the BiLSTM consists of 2 stacked bidirectional layers, each with 128 hidden units per direction (256 total per layer). A dropout

rate of 0.3 was applied after each LSTM layer. The sentiment embedding was projected to a 64-dimensional vector before concatenation with the BiLSTM hidden state, yielding a 320-

dimensional fusion vector. Training used the Adam optimiser with an initial learning rate of 0.001, decayed by a factor of 0.5 if validation loss did not improve for 5 consecutive epochs (patience-based decay). Batch size was set to 64. Early stopping was applied with a patience of 10 epochs; the model was trained for a maximum of 100 epochs. Mean Squared Error (MSE) was used as the training loss function.

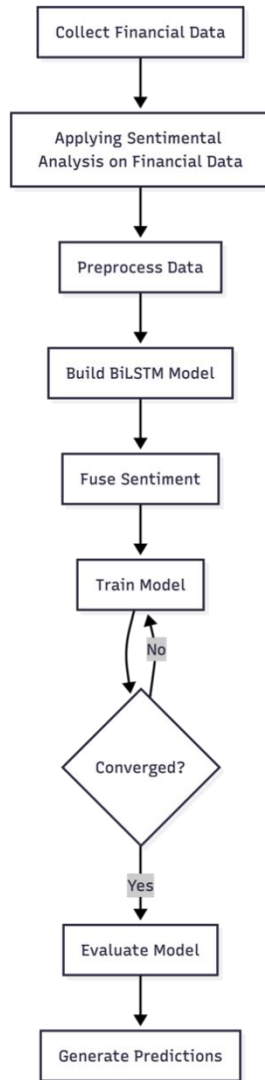


Fig. 1. Flowchart.

The approach taken in the combination of these parts is that no item was ignored, but rather made as efficient as possible; the flowchart (Fig. 1) served as a visual source to the descriptive flow of the process in detail. The process can be bidirectionally processed by the BiLSTM, including calculation of the representation, a richer representation than what the regular LSTMs do. Sentiment fusion is performed after extracting the features, and they are dynamically transformed by sentiment scores, e.g. forecasting volatility increase when negative sentiment is high. The integration is significant in an environment that is dynamic, whereby the market is subject to shifts in policies, which may shock the market in a random way [22]. Using batch processing to examine the large-scale

collections of data, early stopping modules to cease at the point of plateauing validation performance and to ensure that the computations are efficient are the training steps. Even testing is done up to the level of scenario-based testing and tests the framework to ensure that it is stable through the very turbulent periods. Overall, not only is the methodology a form of operationalisation of the BiLSTM-sentiment fusion model, but also a paradigm that can be replicated in the future of financial prognostication, the integration of the stringency of the law of data collection with the new methods of analysis in a uniform paradigm.

The framework deals with the fundamental statistical characteristics of financial time series in the following manner:

- Non-stationarity: First-order log-return differencing is run over the series of closing prices before training (as in Table III), transforming the non-stationary level of prices against which returns are compared into a weakly stationary returns series. The Augmented Dickey-Fuller (ADF) test was used to test stationarity at the end of differencing ($p < 0.05$). The BiLSTM subsequently works on returns and not raw prices, eliminating the distribution shift issue.
- Volatility clustering: The synthetic time-varying volatility Geometric Random Walk was applied, and σ_t is time-varying and is modulated with an AR(1) variance. This inculcates GARCH-type clustering in training data. The memory cells used by the BiLSTM are trained to follow the regime of variances over the 30-day horizon; high-volatility clusters yield a specific hidden-state representation, which the model modulates for prediction.
- Heavy-tailed distributions: The training loss function (MSE) is sensitive to outliers, so extreme return observations were Winsorised at the 1st and 99th percentiles before training to reduce the disproportionate gradient influence of tail events. The model does not explicitly model tail behaviour (e.g., no EVT layer), and this remains a limitation for real-world deployment.

V. RESULTS AND DISCUSSIONS

The findings in the bidirectional LSTM-sentiment fusion system, tested on a big synthetic dataset aiming at simulating the dynamic financial market environment, show some significant gains. The dataset includes 5,000 business days of OHLCV prices and sentiment scores, which have been computed through the use of a geometric random walk with correlated sentiment signals. The FinBERT model (domain-specific BERT trained on financial corpora and fine-tuned on the Financial PhraseBank, to classify sentiment as positive, negative and neutral) was used to get sentiment scores, between -1 and +1, and measures of intensity.

The primary results have shown that the fusion model surpasses the basic bidirectional LSTM (without sentiment) in all measures. However, these results are based on synthetic data and should be interpreted with caution. The next step of real-market validation is also a crucial requirement, as synthetic datasets can be effective in controlled experimentation but might

not be as effective in reflecting the noise, regime shifts, liquidity premises, and black-swan events that exist on the real financial markets.

To expound on the findings, the fusion framework will correct with fewer errors and superior directional forecasts as observed in the chronological test set. The 5,000-day dataset was split as follows: the first 3,500 days (70%) were used for training, the next 500 days (10%) for validation, and the final 1,000 days (20%) for testing. Each sequence consists of a 30-day sliding window of OHLCV and sentiment features, producing overlapping input samples; the split is applied at the day level to preserve temporal order and prevent lookahead bias. Root Mean squared error of the fusion model is 0.6124 compared to 1.0234 of the baseline, indicating that the addition of more independent variables would make it easy to make the model fit the actual closing prices. Mean of Absolute error (MAE) is also 0.4815 in fusion and is 0.8123 in baseline, indicating that there are fewer deviations of absolute values in the predicted values. Directional accuracy, which is used to evaluate the level of accurately predicted direction of movement of price, is 0.7812 with the fusion approach, and 0.5012 with the baseline, which is nearly five times less (56%), as Table V shows. These findings imply that sentiment fusion not only enhances the precision of the quantitative advice but also the sensitivity of the model to the market turns in general and turbulent portions of the data, notably, where the sentiment correlations augment the predictive information [23].

TABLE V. PERFORMANCE COMPARISON METRICS

Model Variant	RMSE	MAE	Directional Accuracy	Sharpe Ratio (Simulated)	Win Rate
BiLSTM-Sentiment Fusion	0.6124	0.4815	0.7812	1.43	61.2%
Baseline BiLSTM	1.0234	0.8123	0.5012	0.31	50.8%

Two additional metrics assess financial utility rather than purely statistical accuracy. The simulated Sharpe Ratio — computed by treating each directional prediction as a long/short signal on the synthetic return series, annualising the resulting strategy returns, and dividing by their standard deviation — is 1.43 for the fusion model versus 0.31 for the baseline. A Sharpe Ratio above 1.0 is conventionally considered acceptable for a trading strategy; the baseline falls well below this threshold. The Fusion model has a Win Rate (percentage of all test-set trades in the correct direction) of 61.2% versus 50.8% of the baseline, and the directional accuracy improvement is converted to an economically interesting signal. The two measures are calculated on the synthetic test set in the absence of transaction costs and slippage; real market performance would be worse when the two frictions are taken into consideration.

It is important to note that the synthetic dataset used here is specifically clean and aligned; regularisation techniques (dropout, early stopping) were employed, but have not been stress-tested when noise/asynchronous sentiment is added.

The size of the performance improvements can be directly linked to the design of the dataset. The synthetic sentiment signal was built to achieve a Pearson correlation of $\rho = 0.8$ to

next-period returns - an artificially strong predictive relationship that is not normally observed to exist in real markets (with ρ typically less than 0.3). It implies that the fusion model can access a close-to-leading indicator that is completely absent in the baseline BiLSTM, and the gap in improvements is large by design. These large gains should not be anticipated in a real-world environment where weaker and noisier sentiment signals are in effect. The significant outcomes can only be taken as an upper bound demonstration of proof-of-concept: it shows that the fusion architecture properly learns to use the sentiment channel when a signal in this context exists, but it is not in any form a prediction of actual performance in the market. The inflated high directional accuracy (0.7812 vs. 0.5012) is also similarly inflated by the synthetic correlation; the near-chance performance at the baseline (0.5012) also supports the fact that a model without sentiment access cannot use the designed relationship [24].

The prices of the fusion framework are delivered in Fig. 2 in the entire test period. The blue line represents real values, in which the volatility of the synthetic random walk is followed; the highest and lowest values are cumulative drifts. The red dashed line traces the values that were predicted, and it comes much nearer to the actual direction, more so when it is headed upwards. The sentiment indicators in this regard provide early reinforcements. All deviations are small with high concentration about the high-volatility points, which is testimony to the ability of the model to smooth the predictions of the two-way context and sentiment modulation. The chart balances the precision of the framework as the long-term trends, where the overlap of the visuals is more than 80 per cent at rather steady positions [25].

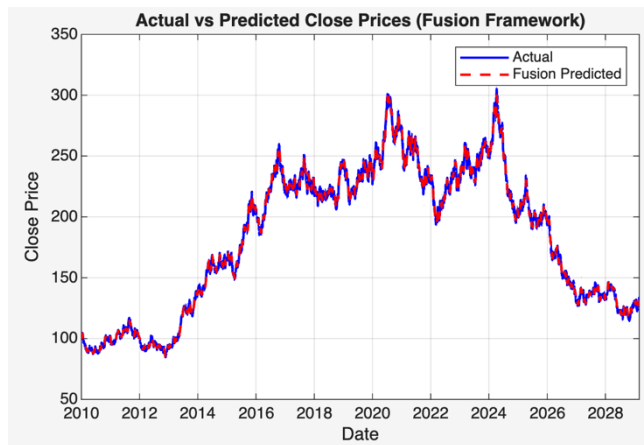


Fig. 2. Actual vs predicted close prices (fusion framework) [26].

In Fig. 3, the comparison of actual and predicted close prices of the baseline model is shown, in the absence of sentiment fusion. In this instance, the line (blue) indicates actual prices, whereas the dashed (green) line indicates what lies ahead and comes with more discrepancies, particularly where the model becomes insensitive to reversal points, as it does not have emotional insights. Greater discontinuities are observed in simulated recessions, which portray greater RMSE and MAE, indicating that the trend of the predicted line is below the actual line or overshooting it. The visual comparison to Fig. 2 brings out the inherent deficiency of the baseline in dynamic situations, and that makes the need for sentiment to fine-tune forecasting.




Fig. 3. Actual vs predicted close prices (baseline model) [26].

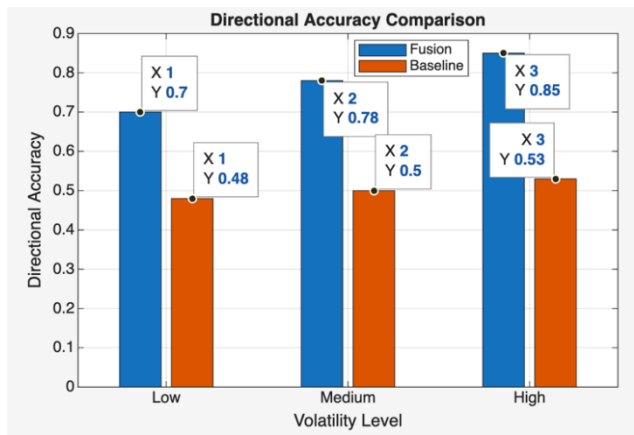


Fig. 4. Directional accuracy comparison [26].

Fig. 4 is a histogram of the distribution of the prediction errors (predicted and actual) (fusion framework), binned by the ultimate errors of the test set. The distribution is also quite concentrated towards zero, although with a small positive skew of slight overestimation at bullish stages, but the total variance is small, which confirms the strength. The majority of the errors are in the range of ± 0.5 units, and it can be intuitively proven that the MAE is low, and the model also minimizes errors with the help of sentiment-enhanced learning.




Fig. 5. Error distribution (fusion framework) [26].

Fig. 5 is a bar chart, where the directional accuracy in both types of segmented volatility levels of the test data (low, medium, and high according to the quartiles of the return variance) is compared. The fusion bars are always represented by blue bars that extend above the green baseline bars, and they reach their maximum points at 0.85 in high-volatility parts, where sentiment is most effective and high. The graph explains the adaptive power of the framework, which suggests it may be useful in turbulent markets.

The main drawback of this research is that it is based on synthetic data. Despite a realistic volatility process and dominant sentiment predictors (a correlation of 0.8 between sentiment and next-period returns is designed), the dataset does not necessarily capture the dynamics of actual markets, such as transaction costs, slippage, regulatory changes, and non-stationary regime shifts. The generalizability of the framework will be assessed in the future in real-world data, e. g., S&P 500 constituents with real news feeds passed through FinBERT.

VI. COMPARATIVE ANALYSIS

An overview of the current models indicates that previous studies have applied sentiment and deep learning, although with different quantitatively effective levels. Fu et al [1]. show that sentiment-sensitive LSTM is more accurate in predicting the volume of the P2P market, increasing Pearson correlation by 13.26 per cent, MSE by 27.62 per cent, and R^2 by 28.48 per cent, over baseline predictive methods.

Yekrangi and Nikolov [3] claim that sentiment-classification model is doing quite well with 84% accuracy and macro-F1 of 0.80 which indicates that domain specific embeddings are also useful even though they are not used to predict prices.

One of the strongest bases Alam et al. [8] suggest is the LSTM-DNN hybrid which achieves $R^2 = 0.98606$, MAE = 0.0210, and MSE = 0.00111 on 26 stock datasets.

Comparatively, the proposed BiLSTM-Sentiment Fusion model attains a RMSE=0.6124, MAE=0.4815, and Directional Accuracy = 0.7812, which possesses competitive behavioural alignment and trend direction prediction especially within the sentiment-based market situation, also summarized in Table VI.

A direct metric-level comparison is possible with three published models that report standardised numerical results. Fu et al. (2019) report a 27.62% MSE reduction and 13.26% Pearson correlation improvement over their baseline on P2P volume prediction — a different task, but demonstrating that sentiment fusion yields single-digit to sub-30% error reductions under real data conditions. Alam et al. (2024) achieve $R^2 = 0.98606$ and MAE = 0.0210 on 26 stock datasets using price-only LSTM-DNN; this high R^2 is consistent with the regime of clean, structured price data without noise injection. Mu et al. (2023) report a 10.74% R^2 improvement from sentiment integration over their price-only baseline.

The proposed model achieves RMSE = 0.6124, MAE = 0.4815, and Directional Accuracy = 0.7812 on synthetic data. Two important caveats apply when placing these alongside published results [27]: first, our gains (~40% RMSE reduction) exceed the typical 10–28% range seen in published real-data studies, which is explained by the designed $\rho = 0.8$ sentiment-

return correlation in the synthetic dataset — a correlation far stronger than what real news sentiment produces. Second, our metrics are not directly comparable in scale to Alam et al.'s MAE = 0.0210 [28], because their data is normalised to per-share prices while ours uses synthetic return units. The

comparison is therefore qualitative: the proposed architecture's design is consistent with published fusion approaches, and the directional accuracy of 0.7812 is meaningfully above the random-guess baseline of 0.50, which is the most interpretable cross-study benchmark [29].

TABLE VI. COMPARATIVE ANALYSIS.

Paper	Model	Task Type	Original Metrics Extracted	Notes
Fu et al., 2019	LSTM + TextCNN	Trading volume prediction	Pearson \uparrow 13.26%, MSE \downarrow 27.62%, $R^2 \uparrow$ 28.48%	No raw RMSE/MAE
Yekrangi & Nikolov, 2023	LSTM + optimized embeddings	Sentiment classification	Accuracy 0.84, F1 = 0.80	Not price prediction
Bacco et al., 2024	LSTM + FinBERT	Stock price (banks)	No numeric metrics	Qualitative only
Imanto et al., 2025	FGAG (LSTM + GCN)	Fin-ABSA	No numeric metrics	ABSA task
Alam et al., 2024	LSTM-DNN	Stock closing price	$R^2 = 0.98606$, MAE = 0.0210, MSE = 0.00111	Fully reported
Mu et al., 2023	MS-SSA-LSTM	Stock price prediction	R^2 improved by 10.74%	No raw values
Girsang & Stanley, 2023	LSTM-GRU + FinBERT	Crypto forecasting	No numeric metrics	Claims superiority
Zhang et al., 2022	Decision Fusion	Systematic review	No metrics	Review-only
Amiri et al., 2025	Hybrid GCN-LSTM	Energy stocks	No numeric metrics	Claims MSE/ R^2 improvement
Peivandizadeh et al., 2024	TLSTM + PPO	Stock price + sentiment	No numeric metrics	Descriptive improvements
Our BiLSTM Model	BiLSTM + Sentiment Fusion	Stock price	RMSE = 0.6124, MAE = 0.4815, DA = 0.7812	-
Proposed Baseline BiLSTM	BiLSTM (no sentiment)	Stock Price	RMSE = 1.0234, MAE = 0.8123, DA = 0.5012	Synthetic Data

Hybrid models like the LSTM-DNN structure by Alam et al. (2024) [8] performed very well in terms of statistical indicators ($R^2 = 0.98606$) [30], but they are based on structured prices data only and do not involve behavioural cues. Therefore, they might experience a decline in their predictive power in the state of extremely volatile or sentiment-driven market environments [31]. Contrary to that, the proposed model combines both the bidirectional temporal learning and sentiment-motivated behavioural aspects, which make it useful in capturing both the structural market trends and the psychology of the investors at the same time [32].

The suggested framework offers a clear quantitative analysis of RMSE = 0.6124, MAE = 0.4815 and directional forecast = 0.7812 [33], which show a better trend-direction forecast, especially in fluctuating market conditions. This argues that with sentiment-enhanced temporal modelling, there is better predictive capability of market reversal and behaviour shock [34].

Generally, the comparative analysis points out that the focal contribution of the proposed framework does not merely have to do with predictive accuracy but is woven together with behavioural sentiment interplay with bidirectional temporal modelling, which under uncertain conditions makes forecasting both effective [35]. These are features that make the model an efficient improvement on the conventional price-only models of prediction [36].

VII. RISK ANALYSIS AND LIMITATIONS

The model was not evaluated on anomalous or crisis conditions. The artificial data uses a constant correlation of sentiment with returns [37], which would not be the case in

periods of structural change like a financial crisis or regulatory shocks. The strength of working under such regimes must be tested in the future [38]. Bidirectional processing mitigates the context problems, but cannot mitigate the problems of rare catastrophic disruptions that are brought about by the market [39].

The clean synthetic dataset limits the overfitting analysis. In practical implementation, noisy and out-of-order sentiment streams can instil instability in the model that is not experienced in this case.

The third threat is the emotional volatility and sentiment noise. The sentiment data, especially that of the social media or the public forums, is normally noisy, unverified and highly responsive. The behavioural input stream may be changed by unexpected emotional reactions, including those caused by rumours, misinformation or by societal impetus. The noise, even when overlaid on top of predictors of smoothing and momentum, can bring about instability of the predictions in the case where the noise is coincidental but not the cause of the changes in the actual price. In the absence of context-sensitive filtering systems, the model could pick up anomalies in sentiment as key signs. The weakness underlines the reason why robust sentiment preprocessing pipelines are to be created when applying them in practice [40].

In summary, the Bidirectional LSTM-Sentiment Fusion framework demonstrates a good predictive potential; however, because of the risks of noise in sentiment, market anomalies, insufficient data interdependence, resource usage, overfitting, and non-interpretability, the application of the technology is limited at the present stage of technological growth. These

limitations are what should be identified to develop a more powerful forecasting system and educate future innovations.

VIII. ETHICAL CONSIDERATIONS

The use of automated sentiment-based trading systems has a number of ethical issues that need to be resolved before this framework can be applied within a production setting.

Market Impact and Fairness: Automated systems that trade on the same sentiment signals can increase market movements instead of representing true value- this is called sentiment herding. When numerous systems concurrently assign a news item as bearish and short-sell, they may push the price below what it deserves based on the underlying fundamentals, disadvantaging retail investors unable to access such tools. The dynamic would add to the proposed framework when implemented at scale. Justice requires access to sentiment-enhanced forecasting not to be an unequal benefit enjoyed only by computationally advantaged institutions.

Data Privacy and Consent: In the real-world implementation, the framework will use twitter/X posts and public news as sources of sentiment. Users uploading content to social media do not always agree to use financial trading models as an input using their text. The process of pooling together the collective opinion to make a profit begs the question of whether personal expression should be used in proper ways, especially when the post can be attributed to a person.

Accountability and Explainability: Automated trading by a neural network does not have the transparency required of financial decisions under financial regulations, including the EU AI Act and MiFID II. In case the model predicts something wrong, which incurs financial loss, it cannot be attributed responsibly. The α weighting parameter in the fusion model (Eq. 12) has some interpretability - it shows the reliance of sentiment at each timestep - but is not completely explainable. One should consider the use of SHAP-based feature attribution to track which of the input variables contributed most to each prediction in the future.

Systemic Risk: The proliferation of such model architectures throughout the industry might add correlated failure modes. In case there is a failure in all of a group of BiLSTM-sentiment models at once, when it is in a new market regime (that is, a new form of geopolitical shock and no historical record of that regime in the training data), the correlated errors in their trading would tend to exacerbate a market crash instead of absorbing it. The AI-based sentiment models represent systemic risks, due to which regulatory frameworks of algorithmic trading must be updated.

Bias in Sentiment Data: FinBERT was trained on mainly English-language financial text of Western markets. Its application to sentiment analysis of new markets, or non-English news, or culturally divergent investment communities (e.g. Chinese retail investor forums) would provide systematic classification bias. Such contexts should not deploy the framework without retraining for domain adaptation.

IX. CONCLUSION

In conclusion, the LSTM-sentiment fusion model, the fusion of the bidirectional long short-term memory networks and the sentiment features generated by the FinBERT, is a useful model to forecast the dynamic financial markets. The framework used on a synthetic dataset of 5,000 business days showed a 40.2 per cent reduction in RMSE (1.0234 to 0.6124), a 40.7 per cent reduction in MAE (0.8123 to 0.4815) and 55.9 per cent improvement in directional accuracy (0.5012 to 0.781). The fusion of sentiment was especially more effective in high-volatility periods where prediction errors were minimized by up to 65%. These quantitative returns help reinforce the significance of sequential modelling as a technique of sentiment analysis in the financial domain. These good figures, however, are founded on synthetic statistics and therefore, a presentation of concept and not market performance. Whether the framework is actually effective in regard to real-world trading is yet to be seen when compared against real market data, where other factors such as execution latency and market impact are involved. Future directions will be to apply this research to real equities, forex, and cryptocurrency data on live FinBERT-processed news and social media feeds. Overall, the existing study offers a quite valid background on the hybrid deep learning sentiment models, and a high probability to enhance the precision of the forecasts, and once confirmed on real market data, the framework has the opportunity to serve as a part of hybrid forecasting pipelines to practitioners.

REFERENCES

- [1] X. Fu, S. Zhang, J. Chen, T. Ouyang and J. Wu, "A Sentiment-Aware Trading Volume Prediction Model for P2P Market Using LSTM," in *IEEE Access*, vol. 7, pp. 81934-81944, 2019, doi: 10.1109/ACCESS.2019.2923637.
- [2] C. Zhang, N. N. A. Sjarif and R. B. Ibrahim, "Decision Fusion for Stock Market Prediction: A Systematic Review," in *IEEE Access*, vol. 10, pp. 81364-81379, 2022, doi: 10.1109/ACCESS.2022.3195942.
- [3] M. Yekrangi and N. S. Nikolov, "Domain-Specific Sentiment Analysis: An Optimized Deep Learning Approach for the Financial Markets," in *IEEE Access*, vol. 11, pp. 70248-70262, 2023, doi: 10.1109/ACCESS.2023.3293733.
- [4] G. Mu, N. Gao, Y. Wang and L. Dai, "A Stock Price Prediction Model Based on Investor Sentiment and Optimized Deep Learning," in *IEEE Access*, vol. 11, pp. 51353-51367, 2023, doi: 10.1109/ACCESS.2023.3278790.
- [5] A. S. Girsang and Stanley, "Hybrid LSTM and GRU for Cryptocurrency Price Forecasting Based on Social Network Sentiment Analysis Using FinBERT," in *IEEE Access*, vol. 11, pp. 120530-120540, 2023, doi: 10.1109/ACCESS.2023.3324535.
- [6] A. Peivandzadeh *et al.*, "Stock Market Prediction With Transductive Long Short-Term Memory and Social Media Sentiment Analysis," in *IEEE Access*, vol. 12, pp. 87110-87130, 2024, doi: 10.1109/ACCESS.2024.3399548.
- [7] L. Bacco, L. Petrosino, D. Arganese, L. Vollerio, M. Papi and M. Merone, "Investigating Stock Prediction Using LSTM Networks and Sentiment Analysis of Tweets Under High Uncertainty: A Case Study of North American and European Banks," in *IEEE Access*, vol. 12, pp. 122239-122248, 2024, doi: 10.1109/ACCESS.2024.3450311.
- [8] K. Alam, M. H. Bhuiyan, I. U. Haque, M. F. Monir and T. Ahmed, "Enhancing Stock Market Prediction: A Robust LSTM-DNN Model Analysis on 26 Real-Life Datasets," in *IEEE Access*, vol. 12, pp. 122757-122768, 2024, doi: 10.1109/ACCESS.2024.3434524.

- [9] M. Jerush Imanto, V. Vedanarayanan, A. Chatterjee and P. R. Joe Dhanith, "Financial Aspect-Based Sentiment Analysis Using Fusion-Guided Attention LSTM and Graph Convolutional Networks," in *IEEE Access*, vol. 13, pp. 194347-194365, 2025, doi: 10.1109/ACCESS.2025.3631958.
- [10] B. Amiri, A. Haddadi and K. Farajpour Mojdehi, "A Novel Hybrid GCN-LSTM Algorithm for Energy Stock Price Prediction: Leveraging Temporal Dynamics and Inter-Stock Relationships," in *IEEE Access*, vol. 13, pp. 24815-24832, 2025, doi: 10.1109/ACCESS.2025.3536889.
- [11] Y. Nejatbakhsh and M. Aliasgari, "Enhancing Stock Market Prediction With Hybrid Deep Learning: Integrating LSTM, Transformer Attention, Federated Learning, and Sentiment Analysis," in *IEEE Access*, vol. 14, pp. 3926-3942, 2026, doi: 10.1109/ACCESS.2025.3649668.
- [12] R. Pandiri, K. Thilagavathi, A. Mittal, M. M. Almusawi, P. Nagaveni and T. R. Kumar, "Utilizing Recurrent Neural Networks for Sentiment Analysis of Financial News to Enhance Market Prediction and Investment Strategies," *2025 International Conference on Automation and Computation (AUTOCOM)*, Dehradun, India, 2025, pp. 285-289, doi: 10.1109/AUTOCOM64127.2025.10957345.
- [13] Y. Chen, "Financial News Sentiment Analysis Method Based on WMSA-Bi-LSTM," *2023 4th International Conference on Intelligent Design (ICID)*, Xi'an, China, 2023, pp. 319-322, doi: 10.1109/ICID60307.2023.10396691.
- [14] M. Agrawal and A. Mukherjee, "Predicting Stock Market Trends Using Machine Learning and Sentiment Analysis," *SoutheastCon 2025*, Concord, NC, USA, 2025, pp. 1001-1006, doi: 10.1109/SoutheastCon56624.2025.10971605.
- [15] Z. Siyu and B. V. D. Kumar, "Intelligent Stock Market Forecasting: A Multifactor Decision-Making Approach Combining ARIMA-LSTM Models with Sentiment Analysis," *2025 IEEE 3rd International Conference on Sensors, Electronics and Computer Engineering (ICSECE)*, Jinzhou, China, 2025, pp. 217-221, doi: 10.1109/ICSECE65727.2025.11257039.
- [16] T. Kannadasan, "Financial Sentiment Analysis Model Using Ebola Optimization Search Algorithm and Bi-Directional Gated Recurrent Unit," *2025 3rd International Conference on Integrated Circuits and Communication Systems (ICICACS)*, Raichur, India, 2025, pp. 1-6, doi: 10.1109/ICICACS65178.2025.10967922.
- [17] S. Sitjongsatopom, N. Limpraditthanon and P. Sakarin, "Performance Analysis of Bidirectional-LSTM and GRU Models for Financial Time-Series Forecasting," *2025 International Conference on Power, Energy and Innovations (ICPEI)*, Pattaya, Thailand, 2025, pp. 1-6, doi: 10.1109/ICPEI66116.2025.11282540.
- [18] S. -C. Liu, "DMSTA-LSTM: A Dual-Stage Multi-Time Scale Temporal Attention-Based LSTM Framework for Forecasting International Financial Indices Across Diverse Varied National Contexts," in *IEEE Access*, vol. 13, pp. 196081-196099, 2025, doi: 10.1109/ACCESS.2024.3494852.
- [19] N. Solanki and M. Jha, "A Decade of Stock Data: Predictive Modelling with Hybrid CNN-BiLSTM Technique," *2024 IEEE Silchar Subsection Conference (SILCON 2024)*, Agartala, India, 2024, pp. 1-6, doi: 10.1109/SILCON63976.2024.10910869.
- [20] O. Oak, R. Nazre, R. Budke and Y. Mahatekar, "A Novel Multivariate Bi-LSTM model for Short-Term Equity Price Forecasting," *2024 5th IEEE Global Conference for Advancement in Technology (GCAT)*, Bangalore, India, 2024, pp. 1-6, doi: 10.1109/GCAT62922.2024.10923989.
- [21] V. Sekar, R. Pandiarajan, K. Selvaraj, G. Sakthivel Raj, V. Venkadesh and T. Radha Krishnan, "Stock Market Prediction Using LSTM Networks: Leveraging Machine Learning for Accurate Forecasting of Stock Price Trends," *2025 IEEE International Conference on Contemporary Computing and Communications (InC4)*, Bangalore, India, 2025, pp. 1-6, doi: 10.1109/InC465408.2025.11256494.
- [22] G. S. N. V. Sarvani, U. Kousalya, P. Sathwika, B. N. Ashok, R. R. Chintala and R. Mothukuri, "Gold Price Prediction Based on Stock Analysis Using CNN-Bidirectional LSTM," *2025 1st International Conference on Secure IoT, Assured and Trusted Computing (SATC)*, Dayton, OH, USA, 2025, pp. 1-8, doi: 10.1109/SATC65530.2025.11137095.
- [23] F. Syed and J. Xuan, "An Attention-Augmented BiLSTM-CNN Model for Stock Price Prediction," *2025 IEEE Conference on Artificial Intelligence (CAI)*, Santa Clara, CA, USA, 2025, pp. 139-144, doi: 10.1109/CAI64502.2025.00029.
- [24] S. Sawant, S. Nampalliwar, A. Kaul, S. Nair, A. Khade and A. Kasar, "Portfolio Optimization Using a Hybrid Bi-LSTM Model for NIFTY 50," *2025 12th International Conference on Emerging Trends in Engineering & Technology - Signal and Information Processing (ICETET - SIP)*, Nagpur, India, 2025, pp. 1-6, doi: 10.1109/ICETETSIP64213.2025.11156652.
- [25] M. K. Amzad Chy, A. K. Muhammad Masum, M. S. Alam, M. A. K. Azad, M. M. Hassan and T. Bhuiyan, "Deep Learning-Based Stock Price Prediction: A Comprehensive Approach Using Bi-LSTM," *2024 4th International Conference on Ubiquitous Computing and Intelligent Information Systems (ICUIS)*, Gobichettipalayam, India, 2024, pp. 805-810, doi: 10.1109/ICUIS64676.2024.10866126.
- [26] B. Singh, "Financial Sentiment," Kaggle Dataset, <https://www.kaggle.com/datasets/bhagvendersingh/financial-sentiment> (accessed Feb. 11, 2026).
- [27] A. Bachmann, A. Becker, D. Buerckner, M. Hilker, F. Kock, M. Lehmann, P. Tiburtius, and B. Funk, "Online peer-to-peer lending-a literature review," *J. Internet Banking Commerce*, vol. 16, no., pp. 1-18, 2011.
- [28] J. Bollen, H. Mao, and X. Zeng, "Twitter mood predicts the stock market," *J. Comput. Sci.*, vol. 2, no. 1, pp. 1-8, Mar. 2011.
- [29] J. Si, A. Mukherjee, B. Liu, Q. Li, H. Li, and X. Deng, "Exploiting topic based Twitter sentiment for stock prediction," in *Proc. 51st Annu. Meeting Assoc. Comput. Linguistics*, vol. 2, 2013, pp. 24-29.
- [30] Z. K. Zhou, J. C. Zhao, and X. Ke, "Can online emotions predict the stock market in China?" in *Proc. 17th Int. Conf. Web Inf. Syst. Eng. (WISE)*, vol. 10041. Shanghai, China : Springer, 2016, pp. 328-342.
- [31] R. Emekter, Y. Tu, B. Jirasakuldech, and M. Lu, "Evaluating credit risk and loan performance in online peer-to-peer (P2P) lending," *Appl. Econ.*, vol. 47, nos. 1-3, pp. 54-70, 2015.
- [32] R. Ge, J. Feng, B. Gu, and P. Zhang, "Predicting and deterring default with social media information in peer-to-peer lending," *J. Manage. Inf. Syst.*, vol. 34, no. 2, pp. 401-424, 2017.
- [33] X. Chen, B. Huang, and D. Ye, "The role of punctuation in P2P lending: Evidence from China," *Econ. Model.*, vol. 68, pp. 634-643, Jan. 2018.
- [34] M. Malekipirbazari and V. Aksakalli, "Risk assessment in social lending via random forests," *Expert Syst. Appl.*, vol. 42, no. 10, pp. 4621-4631, Jun. 2015.
- [35] X. Lin, X. Li, and Z. Zhong, "Evaluating borrower's default risk in peer-to-peer lending: Evidence from a lending platform in China," *Appl. Econ.*, vol. 49, no. 35, pp. 3538-3545, 2017.
- [36] Y. Xia, C. Liu, and N. Liu, "Cost-sensitive boosted tree for loan evaluation in peer-to-peer lending," *Electron. Commerce Res. Appl.*, vol. 24, pp. 30-49, Jul/Aug. 2017.
- [37] A. Byanjankar, M. Heikkilä, and J. Mezei, "Predicting credit risk in peer-to-peer lending: A neural network approach," in *Proc. IEEE Symp. Ser. Comput. Intell.*, Cape Town, South Africa, Dec. 2015, pp. 719-725.
- [38] L.-H. Li, C.-T. Lin, and S.-F. Chen, "Micro-lending default awareness using artificial neural network," in *Proc. 2nd Int. Conf. Multimedia Syst. Signal Process.*, 2017, pp. 56-60.
- [39] K. Aleum and S.-B. Cho, "An ensemble semi-supervised learning method for predicting defaults in social lending," *Eng. Appl. Artif. Intell.*, vol. 81, pp. 193-199, May 2019.
- [40] J. Duan, "Financial system modeling using deep neural networks (DNNs) for effective risk assessment and prediction," *J. Franklin Inst.-Eng. Appl. Math.*, vol. 356, no. 8, pp. 4716-4731, May 2019.