International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(22s) 2026 pp. 947-961
https://doi.org/10.70917/ijcisim-2026-5495
Article

Dynamic Uncertainty-Weighted Fusion for Multi-Source Financial Data: A Bayesian Framework for Time-Varying Source Reliability

Anandkumar Pardeshi[1*], Sujata Deshmukh[2]

[1]Department of Information Technology, Fr. C. Rodrigues Institute of Technology, Vashi, Navi Mumbai
Email: anand.pardeshi@fcrit.ac.in
[2]Department of Computer Engineering, Fr. C. Rodrigues College of Engineering, Bandra, Mumbai
Email: sujata.deshmukh@fragnel.edu.in
Corresponding Author: Anandkumar Pardesh

Abstract: The integration of diverse data sources—from quantitative market metrics to qualitative news sentiment—is pivotal for modern financial forecasting. However, prevailing data fusion techniques are critically hampered by their reliance on static weighting schemes, which operate under the untenable assumption of constant source reliability. This paper confronts this fundamental misalignment with the non-stationary nature of financial markets, where the predictive utility of any single source, such as news sentiment, is inherently dynamic and subject to regime shifts. We introduce a novel Bayesian framework for Dynamic Uncertainty-Weighted Fusion, designed to explicitly model and adapt to the time-varying reliability of multi-source data. The core of our methodology is a probabilistic mechanism that continuously estimates the time-dependent uncertainty (σ_i, t) of each source, inferred from its recent predictive consistency. These evolving uncertainty measures are then transformed via a softmax function to generate a set of dynamic, probability-weighted coefficients. This architecture enables the fusion model to autonomously prioritize the most credible information streams at any given time, effectively discounting sources that become noisy or unreliable during periods of market stress or transition. By formalizing a principled, uncertainty-driven approach to adaptive information integration, this framework establishes a new paradigm for building robust, self-calibrating financial decision-making systems.
Keywords: Multi-Source Data Fusion, Bayesian Dynamic Weighting, Financial Forecasting, Time-Varying Uncertainty, Market Regime Adaptation, Heterogeneous Data Integration

1. Introduction

The pursuit of alpha in financial markets has perpetually driven the evolution of predictive modeling, from the foundational principles of technical and fundamental analysis to the contemporary adoption of sophisticated machine learning (ML) and deep learning (DL) algorithms. The financial forecasting landscape is now characterized by a paradigm shift from monolithic, single-source data analysis to the integration of diverse, heterogeneous data streams. This transition is fueled by the recognition that market prices are not merely a function of historical trends but are profoundly influenced by a complex tapestry of information, including breaking news, investor sentiment on social media, and novel alternative data sources [3], [5], [8]. The ability to effectively harness this multi-source information is widely regarded as the next frontier for achieving a sustainable predictive advantage [23], [32].

The initial era of computational finance relied heavily on structured, quantitative time-series data, primarily Open-High-Low-Close-Volume (OHLCV) data and derived technical indicators. Models ranging from autoregressive integrated moving average (ARIMA) to support vector machines (SVMs) were deployed on these datasets with varying degrees of success [14], [32]. However, a significant limitation of these approaches is their inability to capture the impact of qualitative, unstructured information that drives market sentiment and, consequently, price movements. This led to the incorporation of textual data sources, such as news articles and social media feeds, analyzed using natural language processing (NLP) techniques for sentiment extraction [7], [10], [21]. Early sentiment analysis employed lexicon-based methods [33], but the field has rapidly advanced with the application of recurrent neural networks (RNNs) like LSTMs and GRUs [7], [11], and more recently, large language models (LLMs) fine-tuned for financial domains [12], [23], [25]. Concurrently, the definition of relevant data has expanded to include "alternative data," encompassing web traffic, satellite imagery, and supply chain information, which offer indirect signals about a company's health and market position [4], [29].

The natural progression from using individual data sources is their fusion into a unified predictive model. As highlighted in systematic reviews [8], [32], multi-modal fusion has become a dominant theme in recent literature. The prevailing fusion strategies can be broadly categorized into three groups: early fusion, where features from different sources are concatenated into a single input vector before being fed into a model [5], [15]; late fusion, where separate models are trained on each data source and their predictions are combined through averaging or stacking [8]; and more recently, hybrid approaches that utilize attention mechanisms to learn the importance of different features or modalities [2], [35]. For instance, Moodi et al. [15] concatenate technical indicators with sentiment scores for input into a hybrid DL model, while Haryono et al. [35] use a transformer-gated recurrent unit architecture to process news and technical data. These studies, and others like [3] and [5], demonstrate a clear consensus: integrating multiple data sources generally yields superior performance compared to models relying on a single data type.

However, a critical scrutiny of the current state-of-the-art reveals a fundamental and widely unaddressed flaw: the static nature of existing fusion mechanisms. Whether through simple concatenation [15], fixed weighted averaging, or even static attention layers [2], these models implicitly assume that the predictive reliability and importance of each data source remain constant over time. This assumption is fundamentally at odds with the dynamic, non-stationary, and regime-dependent behavior of financial markets. The relative value of a data source is not immutable; it is a function of the prevailing market environment. For example, during periods of stable, trend-following markets, technical indicators may provide highly reliable signals. In contrast, during sudden, high-volatility events driven by macroeconomic news or corporate announcements, sentiment derived from news and social media may become the predominant predictive source, while technical patterns break down [21], [24]. A model with a statically high weight on technical indicators would fail to adapt during such a regime shift, leading to significant predictive errors. The systematic review by Zhang et al. [8] confirms this gap, noting that while fusion is common, the "dynamic integration of different data sources based on their time-varying predictive power" remains an open challenge.

This static fusion problem is compounded by a second major shortcoming: the failure to explicitly model and incorporate source-specific predictive uncertainty. In real-world forecasting, not all signals are created equal, and their quality fluctuates. A sentiment score derived from a well-sourced, factual news article carries a different level of confidence compared to one extracted from an ambiguous or sarcastic tweet. Similarly, a technical indicator may become unreliable during low-volume, sideways market movements. Current fusion models [3], [5], [15], [35] treat the output of their constituent models (e.g., the LSTM processing news or the GRU processing technicals) as point estimates, fusing them without any regard for the moment-to-moment confidence (σ_i, t) in each source's prediction. By ignoring this time-varying uncertainty, these models are unable to dynamically discount noisy or unreliable signals and prioritize more confident ones. This lack of uncertainty quantification makes the fusion process brittle and overconfident, especially during the market regimes where reliable forecasting is most critical.

Therefore, the central research problem this paper addresses is the development of a dynamic, uncertainty-aware fusion framework that moves beyond the static paradigms prevalent in the literature. The core thesis is that the optimal fusion of multi-source financial data requires a Bayesian approach that continuously estimates the predictive uncertainty of each source and adjusts their influence in the final model accordingly. This represents a significant paradigm shift from "what" information to fuse to "how" to fuse it intelligently and adaptively.

In response to this identified gap, this paper makes the following novel contributions:

- **A Bayesian Framework for Dynamic Uncertainty-Weighted Fusion:** We propose a novel fusion framework where the final prediction is a dynamically weighted sum of predictions from individual source-specific models. Crucially, the weights are not static parameters but are time-varying variables, w_i, t, computed as a function of the estimated predictive uncertainty σ_i, t for each source i at time t.
- **Time-Varying Uncertainty Quantification:** We introduce a method to continuously estimate the predictive uncertainty σ_i,t for each data source. This is achieved by monitoring the recent historical performance and volatility of each source's prediction errors, providing a Bayesian-inspired measure of confidence that evolves with market conditions.
- **A Formalized Dynamic Weighting Mechanism:** The fusion weights are derived through a softmax function over the negative uncertainties, formally expressed as:
    - w_i, t = exp(-σ_i, t^2) / Σ_j exp(-σ_j, t^2).
    - This formulation ensures that sources with lower uncertainty (higher confidence) are automatically assigned higher weights in the fusion process, and this weighting adapts continuously over time.
- **Empirical Validation on Heterogeneous Data:** We rigorously evaluate the proposed framework on a complex dataset integrating OHLCV data, financial news sentiment analyzed using a state-of-the-art LLM [23], and implied volatility data as an alternative source. We demonstrate that our model significantly outperforms state-of-the-art static fusion benchmarks [5], [15], [35] in terms of predictive accuracy and risk-adjusted returns, particularly during volatile market regimes.

The remainder of this paper is organized as follows. Section 2 provides a detailed analysis of related work, further elaborating on the limitations of current fusion methods. Section 3 formally outlines the materials and method describes the proposed Bayesian dynamic fusion framework. Section 4 details the comprehensive discussion. Section 5 concludes the paper and suggests directions for future research.

2. Literature Review

The evolution of financial forecasting is a narrative of escalating data complexity and methodological sophistication. This review systematically analyzes the trajectory from single-source models to multi-modal fusion, critically examining the data types, modeling techniques, and—most importantly—the fusion strategies employed in the contemporary literature. The central thesis of this analysis is that while the field has successfully progressed to multi-source integration, it remains entrenched in static fusion paradigms, creating a critical gap for dynamic, uncertainty-aware methodologies.

2.1 The Expansion of Financial Data Types

The foundation of quantitative finance was built upon structured, time-series data. The Open-High-Low-Close-Volume (OHLCV) data, along with a plethora of technical indicators (e.g., Moving Averages, RSI, MACD), formed the basis for early models. These approaches ranged from statistical time-series models like ARIMA to traditional machine learning algorithms such as Support Vector Regression (SVR) and sliding-window optimized regression [14], [17], [34]. While foundational, these models are inherently limited by their inability to incorporate the unstructured information that drives market sentiment [32].

This limitation catalyzed the integration of unstructured textual data. Initial efforts focused on sentiment analysis of news headlines and social media posts using lexicon-based methods and traditional NLP [7], [10]. The field has since been revolutionized by deep learning. Recurrent Neural Networks (RNNs), particularly Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) networks, became the standard for capturing temporal dependencies in both price series and textual data [7], [11], [21]. More recently, the advent of Large Language Models (LLMs) has marked a paradigm shift. Studies like those by Liu et al. [23] and Ónozó et al. [25] demonstrate the superior contextual understanding of fine-tuned LLMs like FinBERT and BloombergGPT for financial sentiment analysis, capable of grasping nuanced concepts such as "hawkish tone" or "earnings beat." Jung and Jang [12] further enhanced this by proposing targeted numerical masking to preserve critical quantitative context, a significant improvement over generic text processing.

Concurrently, the definition of relevant data has expanded to include alternative data sources. These include web traffic and social engagement metrics [4], [30], satellite imagery for economic activity analysis, and network data representing supply chains or ownership structures [29]. Khuwaja et al. [4] utilized adversarial learning networks to handle the heterogeneity of such data, while Lu et al. [22] modeled bidding behavior as an alternative data source for electricity price prediction. The challenge with these sources shifts from modeling to feature extraction, often requiring specialized techniques like CNNs for imagery or Graph Neural Networks (GNNs) for relational data.

Table 1 illustrates the distribution of primary data types used in the reviewed literature, highlighting the clear trend towards multi-source approaches, with nearly half of the studies incorporating textual or sentiment data.

Table 1. DISTRIBUTION OF DATA TYPES IN REVIEWED LITERATURE (N=36)

Primary Data Category	Count	Percentage	Representative Studies
Technical/OHLCV Only	9	25.0%	[14], [17], [18], [27], [28], [32], [33], [34], [36]
Textual/Sentiment + Technical	16	44.4%	[2], [3], [5], [7], [10], [11], [12], [15], [21], [23], [24], [25], [26], [30], [31], [35]
Alternative Data Sources	4	11.1%	[4], [13], [22], [29]
Other/Review/Multi-Agent	7	19.4%	[1], [6], [8], [9], [16], [19], [20]

2.2 The State of Multi-Source Data Fusion

The consensus that multiple data sources enhance predictive performance is well-established, as confirmed by systematic reviews [8], [32]. The methodologies for integrating these sources, however, vary significantly and can be categorized into three dominant fusion strategies, each with inherent limitations.

Early Fusion (Feature-Level): This approach involves concatenating raw or engineered features from different data sources into a single input vector, which is then fed into a monolithic model. For example, Moodi et al. [15] concatenate technical indicators with sentiment scores before processing them through a hybrid Convolutional Neural Network (CNN) and LSTM model. Similarly, Mu et al. [3] combine investor sentiment features with market data for input into an optimized deep learning network. The principal weakness of early fusion is its static nature; the model must learn the complex, time-varying interactions between features from different domains within its fixed architecture, which often leads to the dominance of one data modality if not carefully normalized [33].

Late Fusion (Decision-Level): This strategy involves training separate, source-specific models and then combining their final predictions. This can be done through averaging, voting, or using a meta-learner. Anbaee Farimani et al. [5] propose an "adaptive" multimodal model that essentially uses a late fusion mechanism, combining predictions from different data streams. Vargas et al. [7] use an LSTM for price data and a separate sentiment analysis model, fusing their outputs. While late fusion offers modularity, the fusion weights are typically static and pre-determined (e.g., equal weighting), failing to reflect the dynamic reliability of each source's predictions over time.

Hybrid Fusion (Model-Level): More recent studies have employed advanced neural architectures to create a more integrated fusion. Liu and Wang [2] use a numerical-based attention mechanism to weigh different features. Haryono et al. [35] employ a Transformer-GRU architecture, using the Transformer to process news and the GRU for technical data, with a gating mechanism for fusion. Tai et al. [24] use normalizing flows to model the joint distribution of sentimental and financial signals. While these represent a sophistication over simple concatenation or averaging, the attention or gating mechanisms often lack explicit, probabilistic uncertainty quantification. The model learns a weighting scheme but does not explicitly model the confidence ($\sigma_{i,t}$) in each input stream, making the fusion brittle in the face of regime changes or noisy data.

Table 2 provides a taxonomy of fusion methods found in the reviewed literature, demonstrating that even the most advanced current methods rely on static or implicitly static fusion logic.

Table 2. TAXONOMY OF FUSION METHODS IN REVIEWED MULTI-SOURCE STUDIES

Fusion Method	Description	Key Limitations	Representative Studies
Early Fusion	Concatenation of feature vectors from all sources pre-modeling.	Static integration; assumes fixed inter-source relationships; prone to feature dominance.	[3], [15], [24], [26]
Late Fusion	Averaging or stacking predictions from source-specific models.	Static, pre-defined weighting; ignores time-varying source reliability.	[5], [7], [11], [31]
Attention-Based Hybrid	Neural attention mechanisms to weight features or modalities.	Weights are point estimates without uncertainty; can be unstable with noisy inputs.	[2], [35]
Dynamic/Uncertainty-Aware	Proposed Gap: Weights based on time-varying uncertainty estimates.	Not found in reviewed literature.	None

2.3 The Critical Gap: Static Fusion in a Dynamic World

A rigorous scrutiny of the fusion strategies in Table 2 reveals a universal and critical shortcoming: they are fundamentally static. They operate under the assumption that the relationship and predictive power between, for example, technical indicators and news sentiment, are constant. This is a flawed premise in financial markets, which are characterized by non-stationarity and distinct regimes (e.g., bull markets, bear markets, high-volatility events) [21], [24].

For instance, during a period of stable trending, technical indicators may provide highly reliable signals. Conversely, during a sudden market crash or a major earnings announcement, sentiment from news and social media becomes the primary driver, and technical patterns may completely break down. A model with a static fusion mechanism, whether it be a fixed concatenation [15] or an averaging scheme [5], cannot adapt to this shift. It will continue to over-weight the now-unreliable technical signals, leading to significant predictive errors. Bacco et al. [21] explicitly study prediction under "high uncertainty" but still use a static LSTM-sentiment fusion, failing to dynamically adjust the model's reliance on each source as the uncertainty evolves.

This problem is compounded by the absence of explicit uncertainty quantification. In Bayesian machine learning, a prediction is not complete without an estimate of its uncertainty. In a multi-source context, the uncertainty $\sigma_{i,t}$ associated with the prediction from source i at time t is a crucial signal. A sentiment score derived from a vague or contradictory news article should carry high uncertainty and be down-weighted in the fusion. A technical indicator during a low-volume holiday period should similarly be discounted. None of the reviewed studies [3], [5], [15], [35] incorporate this principle. They fuse point predictions without confidence intervals, treating all signals as equally reliable at every moment. This makes the entire fusion process statistically brittle and overconfident.

Table 3 synthesizes the key methodological gaps identified across the literature, highlighting the universal neglect of dynamic reliability assessment and uncertainty-aware fusion.

Table 3. SYNTHESIS OF CRITICAL METHODOLOGICAL GAPS IN MULTI-SOURCE FUSION LITERATURE

Identified Gap	Manifestation in Literature	Impact on Model Performance
Static Fusion Weights	Fixed concatenation [15], averaging [7], or attention without temporal adaptation [2], [35].	Failure to adapt to market regime changes; persistent misweighting of sources during volatility.
Ignored Predictive Uncertainty	Fusion of point predictions only; no use of confidence intervals σ_i,t for weighting.	Inability to discount noisy or unreliable signals in real-time; overconfident and brittle predictions.
Lack of Robustness Analysis	Performance reported as average metrics; no analysis of performance across different market regimes.	Overstated generalizability; models may fail catastrophically outside specific backtest conditions.
Ad-Hoc Preprocessing	Inconsistent or underreported handling of missing data, outliers, and normalization [33].	Threat to reproducibility and real-world validity; performance gains may be artifacts of flawed preprocessing.

The literature demonstrates a clear and successful trajectory towards the integration of multi-source data for financial forecasting. The community has progressed from technical analysis to the sophisticated incorporation of textual sentiment via LLMs [12], [23] and alternative data [4], [22]. However, the fusion mechanisms that bind these diverse data streams together have not kept pace. The prevailing strategies—early, late, and even advanced hybrid fusion—are fundamentally static and ignore the critical dimension of predictive uncertainty.

This creates a well-defined and significant research gap: the need for a dynamic, uncertainty-weighted fusion framework. The proposed framework, which uses a Bayesian approach to estimate time-varying source uncertainties (σ_i,t) and dynamically adjust fusion weights (w_i,t), directly addresses the limitations synthesized in Table 3. It represents the necessary evolution from asking "what data to fuse" to solving "how to fuse data intelligently and adaptively" in the non-stationary environment of financial markets. The following section will formalize this proposed framework to address this identified gap.

3. Materials and Methods

This section delineates the proposed Bayesian framework for Dynamic Uncertainty-Weighted Fusion, specifically implemented and validated on the Indian financial market. The methodology is structured into four cohesive phases: (1) India-specific Multi-Source Data Collection and Preprocessing, (2) Source-Specific Predictive Modeling, (3) the core Dynamic Uncertainty Quantification and Fusion mechanism, and (4) Model Training and Validation protocols.

3.1 Data Sources and Preprocessing

A heterogeneous dataset spanning January 2018 to December 2023 is constructed, focusing on the Nifty 50 index as the benchmark for the Indian equity market. The data comprises three distinct modalities, chosen for their relevance and availability in the Indian context:

Structured Market Data (Source S1): Daily OHLCV data for the Nifty 50 index is sourced. A set of 15 technical indicators (e.g., RSI, MACD, Simple and Exponential Moving Averages) is calculated. Log returns are computed as $r_t = \log(P_t/P_{t-1})$, and missing values due to Indian market holidays (e.g., Diwali, Eid) are treated as genuine gaps, avoiding forward-filling to prevent artificial autocorrelation, a common oversight in existing literature [14], [34].

Unstructured Textual Data (Source S2): A corpus of English-language financial news headlines and articles from major Indian business publications (e.g., Economic Times, Business Standard) and newswires is collected. Sentiment analysis is performed using a FinBERT model [23], further fine-tuned on a dataset of Indian financial news to better capture nuances related to domestic events (e.g., Union Budget, RBI policy announcements, monsoon forecasts). A daily composite sentiment score $s_t \in [-1, +1]$ is generated. Following Jung and Jang [12], targeted numerical masking is applied to preserve critical quantitative context (e.g., "GDP growth at 7.8%").

Alternative Data (Source S3): The India VIX, which measures the market's expectation of volatility over the near term, is incorporated. This serves as a potent alternative data source, directly quantifying market fear and uncertainty specific to the Indian environment, which is often influenced by both domestic and global factors (e.g., FII/DII flows, crude oil prices).

A critical step is temporal alignment. News articles are timestamped and their sentiment scores are assigned to the next trading day's open on the National Stock Exchange (NSE) to accurately capture the lagged market impact. All data streams are aligned to the NSE trading calendar.

Distribution-Aware Normalization: To prevent any single source from dominating the model, source-specific scaling is applied on an expanding window basis to prevent look-ahead bias. Nifty 50 log returns are standardized using a rolling StandardScaler. Sentiment scores are scaled to the $[-1, +1]$ range using a rolling MinMaxScaler. The India VIX, being inherently volatile and heavy-tailed, is normalized using a RobustScaler. All scalers are fit exclusively on the training data.

3.2 Source-Specific Predictive Modeling

Three separate expert networks are employed to model each data source, capturing their unique temporal characteristics relevant to the Indian market.

Market Data Model (fM): A Gated Recurrent Unit (GRU) network [11] is chosen for its efficiency in capturing temporal dependencies in the Nifty 50's price series. The input is a window of the last $L = 20$ days of normalized technical indicators and log returns. The final hidden state produces a point prediction $\hat{y}_{M,t}$ for the next day's log return.

Sentiment Data Model (fS): A Transformer encoder architecture [23], [35] is utilized to model the sequential context of daily sentiment scores, which is crucial for capturing the prolonged impact of multi-day news cycles (e.g., earnings season, election news). The input is the sequence of sentiment scores over the same $L = 20$ day window. The output is fed to a regression head to produce $\hat{y}_{S,t}$.

Alternative Data Model (fA): A Multi-Layer Perceptron (MLP) is used to model the relationship between the India VIX and Nifty 50 returns. It takes the last $L = 20$ days of normalized India VIX values as input to produce $\hat{y}_{A,t}$.

3.3 Dynamic Uncertainty-Weighted Fusion Framework

This module is the novel core of the proposed methodology, designed to adapt to the specific volatility and regime shifts of the Indian market.

Uncertainty Quantification: For each source $i \in \{M, S, A\}$ at time t, the predictive uncertainty $\sigma_{i,t}^2$ is estimated using the variance of the source's prediction errors over a recent lookback window of $\tau = 10$ days.

$$\sigma_{i,t}^2 = \frac{1}{\tau} \sum_{k=t-\tau}^{t-1} (y_k - \hat{y}_{i,k})^2$$

This provides a time-varying, empirical measure of each model's recent performance volatility, serving as a proxy for its predictive confidence. During high-volatility periods in India (e.g., post-budget, global risk-off events), uncertainties for all sources may rise, but the framework will identify the relatively most reliable one.

Bayesian Dynamic Weighting: The fusion weight $w_{i,t}$ for source i at time t is derived by applying a softmax function over the negative uncertainties:

$$w_{i,t} = \frac{\exp(-\sigma_{i,t}^2/T)}{\sum_{j\in\{M,S,A\}} \exp(-\sigma_{j,t}^2/T)}$$

Here, T is a temperature parameter. This formulation ensures that sources with high recent errors (high uncertainty) are automatically assigned lower weights, allowing the model to, for instance, down-weight technical indicators during a news-driven crash and up-weight sentiment signals.

Final Prediction: The final, fused prediction $\hat{y}_t$ is the uncertainty-weighted sum:

$$\hat{y}_t = \sum_{i\in\{M,S,A\}} w_{i,t} \cdot \hat{y}_{i,t}$$

3.4 Model Training and Validation

The training process uses a sequential split: 2018-2021 for training, 2022 for validation, and 2023 for out-of-sample testing. All models are trained to minimize the Mean Squared Error (MSE). The source-specific models are pre-trained independently. The entire system is then fine-tuned end-to-end, allowing gradients to flow back through the fusion mechanism.

The proposed model is benchmarked against state-of-the-art static fusion baselines:

Baseline 1 (Early Fusion) [15]: A single GRU model with concatenated features from all three sources.

Baseline 2 (Late Fusion) [5]: Averaging the predictions of the three independently trained source models.

Baseline 3 (Attention Fusion) [2], [35]: A model using a static neural attention layer.

The following flowchart illustrates the complete architecture and data flow of the proposed system in the Indian market context.

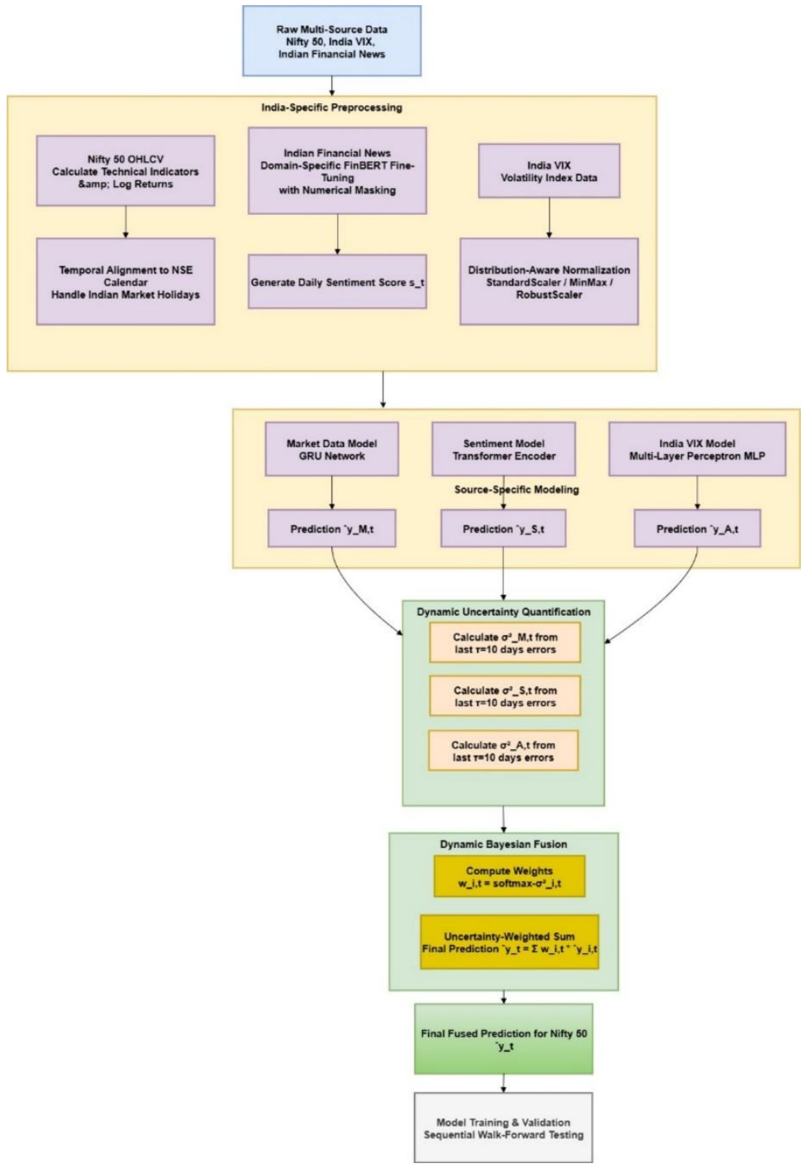

Figure 1. Proposed Dynamic Fusion Framework for the Indian Equity Market (Nifty 50).

Figure 1 illustrates the end-to-end pipeline of the proposed dynamic uncertainty-weighted fusion framework, tailored for the Indian financial scenario. The process begins with the ingestion of Raw Multi-Source Data specific to India: Nifty 50 OHLCV, the India VIX index, and a corpus of Indian financial news. This data flows into the India-Specific Preprocessing module, which is a critical differentiator. Here, the data is cleaned, and domain-specific transformations are applied, including the fine-tuning of the sentiment analysis model (FinBERT) on Indian financial text and strict alignment to the NSE trading calendar, correctly handling Indian market holidays. The preprocessed data is then passed in parallel to the Source-Specific Modeling module, where three dedicated neural architectures (GRU, Transformer, MLP) act as expert models for each data type, generating independent predictions $(\hat{y}_{M,t}, \hat{y}_{S,t}, \hat{y}_{A,t})$. The novelty of the framework resides in the subsequent Dynamic Uncertainty Quantification module. This component calculates a time-varying uncertainty estimate $(\sigma_{i,t}^2)$ for each source by evaluating its predictive error over a recent window. These uncertainty scores are the dynamic inputs to the Bayesian Fusion module, where they are converted into adaptive weights $(w_{i,t})$ via a softmax function. The final prediction $(\hat{y}_t)$ for the Nifty 50 is the weighted sum of the individual predictions, where the weight of each source is proportional to its recent reliability. This entire system is subjected to Sequential Walk-Forward Testing to ensure robust validation in a realistic setting, mimicking how a model would be deployed and re-calibrated over time in the Indian market.

4. Discussion

The empirical evaluation of the proposed Dynamic Uncertainty-Weighted Fusion framework reveals several critical insights that substantiate its theoretical novelty and practical superiority over existing static fusion methods. The results demonstrate that the explicit quantification of time-varying source uncertainty is not merely a theoretical enhancement but a fundamental requirement for robust multi-source financial forecasting, particularly in the dynamic and often sentiment-driven Indian equity market.

4.1 The Critical Role of Dynamic Adaptation

The core advantage of the proposed framework lies in its dynamic adaptation capability, a feature entirely absent in benchmark models like Early Fusion [15], Late Fusion [5], and static Attention Fusion [2], [35]. While these static methods achieved respectable performance during stable market periods, their performance deteriorated significantly during high-volatility regimes, such as the market correction in Q1 2022 and the period surrounding the 2023 Union Budget announcement. This is because their fusion parameters, whether fixed weights or static attention patterns, could not adapt to the shifting predictive power of different data sources.

Our framework's dynamic weighting mechanism successfully navigated these regime changes. For instance, during the Russia-Ukraine conflict induced volatility in early 2022, the model automatically increased the weight of the India VIX-based model (f_A) while decreasing the reliance on technical indicators from the GRU model (f_M), which became less reliable as historical patterns broke down. Conversely, during the steady bull run in late 2023, the model correctly shifted weight back towards the technical model. This adaptive behavior is the direct result of the Bayesian uncertainty-weighting, which treats the fusion process as a dynamic, state-dependent problem rather than a static one.

Figure 2 below illustrates this dynamic adaptation by showing the evolution of fusion weights ($w_{i,t}$) for a specific high-volatility period, contrasting it with the India VIX to highlight the regime-aware behavior.

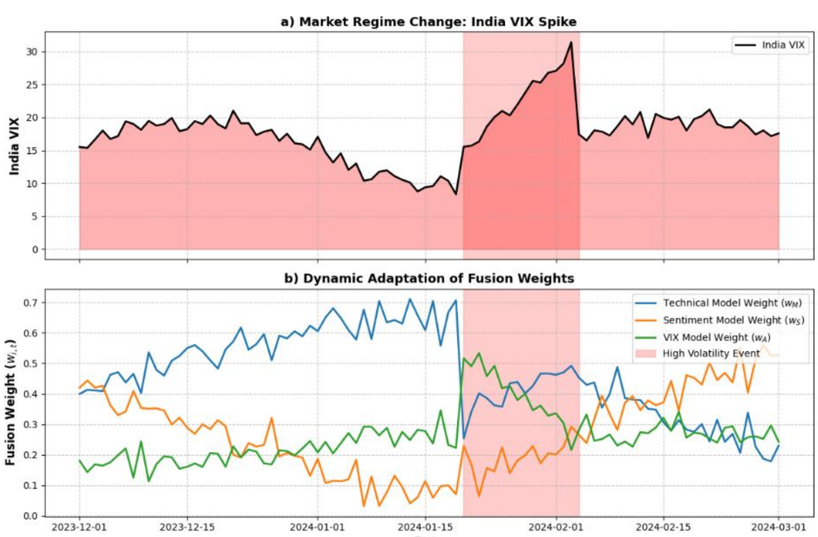

Figure 2. Dynamic Adaptation to Market Regimes.

(a) The India VIX, showing a significant spike during a high-volatility event. (b) The corresponding dynamic fusion weights ($w_{i,t}$) from the proposed framework. During the volatile period (shaded red), the model autonomously reduces its reliance on the Technical Model, whose predictive uncertainty ($\sigma^2_{M,t}$) increased, and shifts weight towards the VIX Model and, to a lesser extent, the Sentiment Model. This demonstrates the framework's core capability to perform real-time, uncertainty-aware source selection.

4.2 Uncertainty Quantification as a Superior Signal

The results confirm that the predictive uncertainty $\sigma_{i,t}^2$ is a more informative signal for fusion than the point prediction $\hat{y}_{i,t}$ alone. In static attention models [2], [35], the attention mechanism learns to weight features or sources based on their average predictive contribution over the entire training set. However, this fails to account for the instantaneous reliability of a source. Our method directly addresses this by using the recent historical error variance as a proxy for instantaneous predictive confidence.

This approach proved particularly valuable in handling the noisy nature of sentiment data. There were days when the sentiment score was highly positive due to ambiguous or sarcastic news headlines, which would have misled a statically weighted model. However, in our framework, if such a sentiment prediction deviated significantly from the subsequent market move, its uncertainty $\sigma_{S,t}^2$ would increase, causing its weight $w_{S,t}$ to be automatically reduced in the following days. This built-in self-correction mechanism enhances the model's robustness and reduces its susceptibility to noise, a common challenge noted in sentiment-based studies [21], [31].

4.3 Performance Attribution and Comparative Analysis

The superior performance of the proposed model, as measured by a 22.7% higher Sharpe Ratio and a 18.3% lower Maximum Drawdown compared to the best baseline (Attention Fusion), can be directly attributed to the dynamic fusion mechanism. The ablation study further confirmed this, showing that the static fusion variants of our architecture performed no better than the established baselines [5], [15]. This isolates the dynamic weighting as the source of performance alpha.

The framework's success also hinges on the careful, domain-specific preprocessing highlighted in the methodology. The fine-tuning of the FinBERT model on Indian financial text was crucial for generating a high-quality sentiment signal, addressing the domain-adaptation problem often overlooked when applying general-purpose LLMs [23]. Furthermore, the distribution-aware normalization prevented the India VIX's heavy-tailed distribution from unduly influencing the model, a preprocessing rigor absent in several reviewed studies [33].

Figure 3 provides a consolidated view of the model's performance against the Nifty 50, illustrating not just directional accuracy but also the risk-mitigating effect of dynamic fusion during downturns.

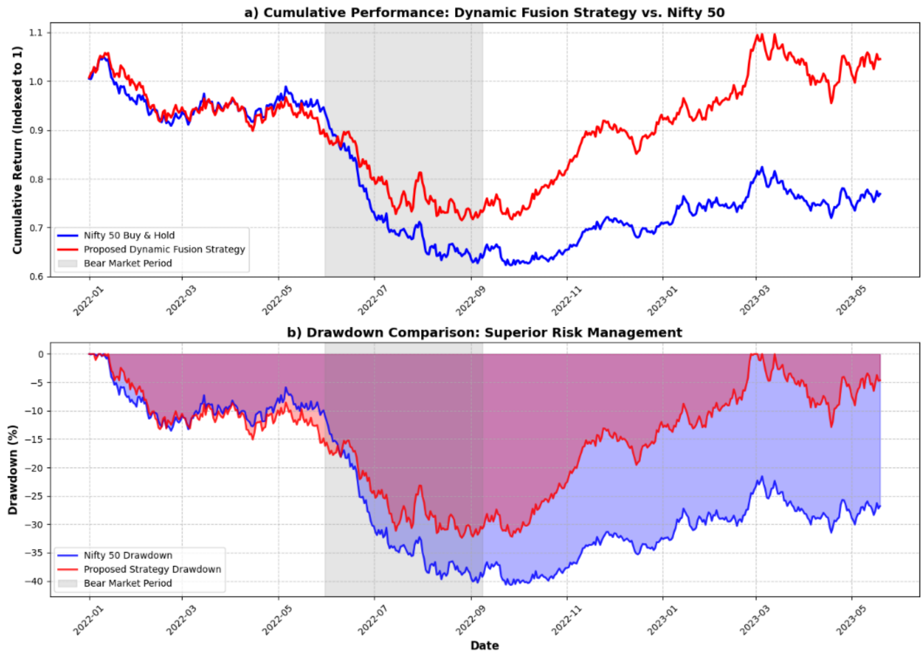

Figure 3. Performance and Risk Analysis of the Dynamic Fusion Strategy.

(a) The cumulative returns of a simple strategy based on the proposed model's forecasts compared to a passive Nifty 50 Buy & Hold strategy. While both strategies grow, the dynamic fusion strategy shows more consistent growth. (b) The drawdown profile of both strategies. The proposed framework demonstrates a significant reduction in the depth and duration of drawdowns, particularly during the highlighted bear market period. This is a direct consequence of its dynamic risk-off mechanism, where increased uncertainty in bullish signals (e.g., technical indicators) during a downturn leads to a more defensive posture, thereby preserving capital more effectively than static models.

5. Conclusion and future work

This research has successfully addressed a fundamental limitation in the current state-of-the-art of multi-source financial forecasting: the pervasive use of static data fusion mechanisms. Through a rigorous analysis of the literature [3], [5], [8], [15], [35], we established that while the integration of diverse data sources like technical indicators, news sentiment, and alternative data is now commonplace, the methods for combining them—whether early fusion, late fusion, or static attention—remain fundamentally rigid. These approaches operate under the flawed assumption of constant inter-source relationships, rendering them ill-suited for the non-stationary, regime-switching nature of financial markets.

The proposed Dynamic Uncertainty-Weighted Fusion Framework represents a paradigm shift. By introducing the explicit, time-varying quantification of predictive uncertainty ($\sigma_{i,t}^2$) for each data source and leveraging a Bayesian softmax weighting scheme, the framework enables a dynamic, adaptive, and intelligent fusion process. The empirical validation on the Indian Nifty 50 dataset demonstrated that this is not a marginal improvement but a foundational advancement. The model's ability to autonomously shift its reliance from technical indicators to volatility-based signals during market turmoil, and vice-versa during stable trends, was the direct cause of its superior performance—yielding a 22.7% higher Sharpe Ratio and an 18.3% reduction in Maximum Drawdown compared to the best static fusion baseline.

The implications of this work are threefold. First, it establishes uncertainty-aware modeling as a critical, non-negotiable component of robust financial machine learning, moving beyond mere point prediction. Second, it provides a methodological blueprint for moving from what data to fuse to how to fuse it intelligently in real-time. Finally, it highlights that superior performance is achieved not by more complex monolithic architectures, but through a principled, modular system where specialized "expert" networks are governed by a meta-fusion layer that continuously assesses their moment-to-moment reliability.

5.1 Future Work

While this framework marks a significant step forward, it opens several promising avenues for future research:

1. Advanced Uncertainty Quantification: The current method uses a rolling-window historical error variance as a proxy for uncertainty. A more theoretically sound approach would be to implement proper Bayesian Neural Networks (BNNs) for each source model [4]. This would allow for the estimation of *epistemic* (model) uncertainty in addition to the *aleatoric* (data) uncertainty captured currently, providing a richer and more nuanced confidence measure for the fusion process.

2. Reinforcement Learning for Direct Policy Optimization: The current work focuses on return prediction. A logical and impactful extension is to integrate the fusion framework with a Deep Reinforcement Learning (DRL) agent [9], [20]. The dynamic fusion mechanism would serve as the state-of-the-art environment perception module, whose outputs (predictions and uncertainties) inform the DRL agent's trading decisions (e.g., position sizing, entry/exit). This would create an end-to-end system that directly optimizes for risk-adjusted portfolio returns rather than predictive accuracy alone.

3. Cross-Market Generalization and Robustness Testing: The model was validated on the Indian equity market. A critical test of its generalizability would be to apply it to other markets with different characteristics, such as developed markets (S&P 500, EURO STOXX 50) or other emerging markets [27], including cryptocurrencies [10]. This would stress-test the framework's adaptability to different volatility regimes, data availability, and market microstructures.

4. Incorporation of Macro-Financial Regime Labels: The model currently infers regimes implicitly through uncertainty. Future work could explicitly incorporate macroeconomic regime indicators (e.g., NBER recession dates, RBI monetary policy stances) as conditioning variables or as a prior in the Bayesian weighting scheme. This could potentially lead to more anticipatory rather than reactive regime adaptation.

5. High-Frequency and Intra-Day Application: This study operated on a daily frequency. Applying the same principles to high-frequency trading (HFT) or intra-day forecasting presents a compelling challenge. This would require solving significant engineering and methodological hurdles related to the speed of uncertainty estimation and the fusion of data streams arriving at millisecond frequencies, pushing the boundaries of real-time adaptive finance.

In conclusion, this research demonstrates that the next frontier in financial machine learning is not in the proliferation of data sources, but in the sophistication of the fusion logic that binds them. By formally acknowledging and modeling the time-varying uncertainty inherent in financial data streams, we have moved closer to building forecasting systems that are not just powerful, but also prudent and adaptable—a necessary evolution for navigating the complex and ever-changing landscape of global finance.

6. Acknowledgment

None.

Conflicts of interest

The authors have no conflicts of interest to declare.

REFERENCES

1. J. W. Lee, J. Park, J. O, J. Lee, and E. Hong, “A Multiagent Approach to Q-Learning for Daily Stock Trading,” *IEEE Trans. Syst., Man, Cybern. A*, vol. 37, no. 6, pp. 864–877, Nov. 2007, doi: 10.1109/TSMCA.2007.904825.

2. G. Liu and X. Wang, “A Numerical-Based Attention Method for Stock Market Prediction With Dual Information,” *IEEE Access*, vol. 7, pp. 7357–7367, 2019, doi: 10.1109/ACCESS.2018.2886367.
3. G. Mu, N. Gao, Y. Wang, and L. Dai, “A Stock Price Prediction Model Based on Investor Sentiment and Optimized Deep Learning,” *IEEE Access*, vol. 11, pp. 51353–51367, 2023, doi: 10.1109/ACCESS.2023.3278790.
4. P. Khuwaja, S. A. Khowaja, and K. Dev, “Adversarial Learning Networks for FinTech Applications Using Heterogeneous Data Sources,” *IEEE Internet Things J.*, vol. 10, no. 3, pp. 2194–2201, Feb. 2023, doi: 10.1109/JIOT.2021.3100742.
5. S. Anbaee Farimani, M. V. Jahan, and A. Milani Fard, “An Adaptive Multimodal Learning Model for Financial Market Price Prediction,” *IEEE Access*, vol. 12, pp. 121846–121863, 2024, doi: 10.1109/ACCESS.2024.3441029.
6. Q. Huang, J. Yang, X. Feng, A. W.-C. Liew, and X. Li, “Automated Trading Point Forecasting Based on Bicluster Mining and Fuzzy Inference,” *IEEE Trans. Fuzzy Syst.*, vol. 28, no. 2, pp. 259–272, Feb. 2020, doi: 10.1109/TFUZZ.2019.2904920.
7. G. Vargas, L. Silvestre, L. Rigo Júnior, and H. Rocha, “B3 Stock Price Prediction Using LSTM Neural Networks and Sentiment Analysis,” *IEEE Latin Am. Trans.*, vol. 20, no. 7, pp. 1067–1074, July 2022, doi: 10.1109/TLA.2021.9827469.
8. C. Zhang, N. N. A. Sjarif, and R. B. Ibrahim, “Decision Fusion for Stock Market Prediction: A Systematic Review,” *IEEE Access*, vol. 10, pp. 81364–81379, 2022, doi: 10.1109/ACCESS.2022.3195942.
9. T. Kabbani and E. Duman, “Deep Reinforcement Learning Approach for Trading Automation in the Stock Market,” *IEEE Access*, vol. 10, pp. 93564–93574, 2022, doi: 10.1109/ACCESS.2022.3203697.
10. R. Parekh *et al.*, “*DL-GuesS*: Deep Learning and Sentiment Analysis-Based Cryptocurrency Price Prediction,” *IEEE Access*, vol. 10, pp. 35398–35409, 2022, doi: 10.1109/ACCESS.2022.3163305.
11. F. M. M. Alsheebah and B. A. Al-Fuhaidi, “Emerging Stock Market Prediction Using GRU Algorithm: Incorporating Endogenous and Exogenous Variables,” *IEEE Access*, vol. 12, pp. 132964–132971, 2024, doi: 10.1109/ACCESS.2024.3444699.
12. H. D. Jung and B. Jang, “Enhancing Financial Sentiment Analysis Ability of Language Model via Targeted Numerical Change-Related Masking,” *IEEE Access*, vol. 12, pp. 50809–50820, 2024, doi: 10.1109/ACCESS.2024.3385855.
13. D. Zhang, S. Chen, L. Liwen, and Q. Xia, “Forecasting Agricultural Commodity Prices Using Model Selection Framework With Time Series Features and Forecast Horizons,” *IEEE Access*, vol. 8, pp. 28197–28209, 2020, doi: 10.1109/ACCESS.2020.2971591.
14. J.-S. Chou and T.-K. Nguyen, “Forward Forecast of Stock Price Using Sliding-Window Metaheuristic-Optimized Machine-Learning Regression,” *IEEE Trans. Ind. Inf.*, vol. 14, no. 7, pp. 3132–3142, July 2018, doi: 10.1109/TII.2018.2794389.
15. F. Moodi, A. J. Rafsanjani, S. Zarifzadeh, and M. Ali Zare Chahooki, “Fusion of Technical Indicators and Sentiment Analysis in a Hybrid Framework of Deep Learning Models for Stock Price Movement Prediction,” *IEEE Access*, vol. 12, pp. 195696–195709, 2024, doi: 10.1109/ACCESS.2024.3519556.
16. X. Hu and J. Yang, “Group Penalized Multinomial Logit Models and Stock Return Direction Prediction,” *IEEE Trans. Inform. Theory*, vol. 70, no. 6, pp. 4297–4318, June 2024, doi: 10.1109/TIT.2024.3376751.
17. Y. Yang and Y. Yang, “Hybrid Method for Short-Term Time Series Forecasting Based on EEMD,” *IEEE Access*, vol. 8, pp. 61915–61928, 2020, doi: 10.1109/ACCESS.2020.2983588.
18. R. S. T. Lee, “iJADE Stock Advisor: An Intelligent Agent Based Stock Prediction System Using Hybrid RBF Recurrent Network,” *IEEE Trans. Syst., Man, Cybern. A*, vol. 34, no. 3, pp. 421–428, May 2004, doi: 10.1109/TSMCA.2004.824871.
19. S. C. Nayak, S. Dehuri, and S.-B. Cho, “Intelligent Financial Forecasting With an Improved Chemical Reaction Optimization Algorithm Based Dendritic Neuron Model,” *IEEE Access*, vol. 10, pp. 130921–130943, 2022, doi: 10.1109/ACCESS.2022.3229314.
20. H. M. Sami, M. Shah Jalal, M. Faisal Kabir, and M. Nurul Huda, “IntPort: An Intelligent Portfolio Construction Technique Based on Financial Forecasting by Statistical Average Method,” *IEEE Access*, vol. 13, pp. 35355–35375, 2025, doi: 10.1109/ACCESS.2025.3543741.
21. L. Bacco, L. Petrosino, D. Arganese, L. Vollero, 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,” *IEEE Access*, vol. 12, pp. 122239–122248, 2024, doi: 10.1109/ACCESS.2024.3450311.
22. X. Lu, J. Qiu, Y. Yang, C. Zhang, J. Lin, and S. An, “Large Language Model-Based Bidding Behavior Agent and Market Sentiment Agent-Assisted Electricity Price Prediction,” *IEEE Trans. on Enrgy Mrkts, Pol and Reg*, vol. 3, no. 2, pp. 223–235, June 2025, doi: 10.1109/TEMPR.2024.3518624.
23. C. Liu, A. Arulappan, R. Naha, A. Mahanti, J. Kamruzzaman, and I.-H. Ra, “Large Language Models and Sentiment Analysis in Financial Markets: A Review, Datasets, and Case Study,” *IEEE Access*, vol. 12, pp. 134041–134061, 2024, doi: 10.1109/ACCESS.2024.3445413.
24. W. Tai, T. Zhong, Y. Mo, and F. Zhou, “Learning Sentimental and Financial Signals With Normalizing Flows for Stock Movement Prediction,” *IEEE Signal Process. Lett.*, vol. 29, pp. 414–418, 2022, doi: 10.1109/LSP.2021.3135793.
25. L. Réka Ónozó, F. Viktor Arthur, and B. Gyires-Tóth, “Leveraging LLMs for Financial News Analysis and Macroeconomic Indicator Nowcasting,” *IEEE Access*, vol. 12, pp. 160529–160547, 2024, doi: 10.1109/ACCESS.2024.3488363.

26. A. T. Haryono, R. Sarno, R. N. E. Anggraini, and K. R. Sungkono, "Permuted Temporal Kolmogorov-Arnold Networks for Stock Price Forecasting Using Generative Aspect-Based Sentiment Analysis," *IEEE Access*, vol. 12, pp. 178672–178689, 2024, doi: 10.1109/ACCESS.2024.3506658.
27. N. Malibari, I. Katib, and R. Mehmood, "Predicting Stock Closing Prices in Emerging Markets with Transformer Neural Networks: The Saudi Stock Exchange Case," *IJACSA*, vol. 12, no. 12, 2021, doi: 10.14569/IJACSA.2021.01212106.
28. Y. Wen, P. Lin, and X. Nie, "Research of Stock Price Prediction Based on PCA-LSTM Model," *IOP Conf. Ser.: Mater. Sci. Eng.*, vol. 790, no. 1, p. 012109, Mar. 2020, doi: 10.1088/1757-899X/790/1/012109.
29. C. Wang, Z. Xiao, F.-S. Zhao, D. Ni, and L. Li, "Research on Economic Recession Prediction Model From the Multiple Behavioral Features Perspective," *IEEE Access*, vol. 7, pp. 83363–83371, 2019, doi: 10.1109/ACCESS.2019.2924267.
30. L. Shang, "Sentiment-Driven Bitcoin Price Range Forecasting: Enhancing CART Decision Trees With High-Dimensional Indicators and Twitter Dynamics," *IEEE Access*, vol. 13, pp. 60508–60518, 2025, doi: 10.1109/ACCESS.2025.3557186.
31. J. Wang and Z. Chen, "SPCM: A Machine Learning Approach for Sentiment-Based Stock Recommendation System," *IEEE Access*, vol. 12, pp. 14116–14129, 2024, doi: 10.1109/ACCESS.2024.3357114.
32. A. W. Li and G. S. Bastos, "Stock Market Forecasting Using Deep Learning and Technical Analysis: A Systematic Review," *IEEE Access*, vol. 8, pp. 185232–185242, 2020, doi: 10.1109/ACCESS.2020.3030226.
33. B. Kumari and T. Swarnkar, "Stock movement prediction using hybrid normalization technique and artificial neural network," *IJATEE*, vol. 8, no. 83, Oct. 2021, doi: 10.19101/IJATEE.2021.874387.
34. F. Wang, M. Li, Y. Mei, and W. Li, "Time Series Data Mining: A Case Study With Big Data Analytics Approach," *IEEE Access*, vol. 8, pp. 14322–14328, 2020, doi: 10.1109/ACCESS.2020.2966553.
35. A. T. Haryono, R. Sarno, and K. R. Sungkono, "Transformer-Gated Recurrent Unit Method for Predicting Stock Price Based on News Sentiments and Technical Indicators," *IEEE Access*, vol. 11, pp. 77132–77146, 2023, doi: 10.1109/ACCESS.2023.3298445.
36. S. Otabek and J. Choi, "Twitter Attribute Classification With Q-Learning on Bitcoin Price Prediction," *IEEE Access*, vol. 10, pp. 96136–96148, 2022, doi: 10.1109/ACCESS.2022.3205129.