Grenze International Journal of Engineering and Technology, June Issue

Stock Market Prediction with Machine Learning: A Regime-Aware Hybrid Transformer Framework

Monika Kamble[1], Aashlesh Wawge[2], Shravani Nomulwar[3], Aditi Sakhalkar[4] and Emmanuel Mark[5]
[1-5]Department of Information Technology, Pune Institute of Computer Technology, Pune, India
Email: monikakamble800@gmail.com, aashleshvwawge@gmail.com, shravaninomulwar@gmail.com, aditisak19@gmail.com, emmanuelm@pict.edu

***Abstract*—Stock market prediction is uniquely challenging due to the interplay of price trends, fundamentals, sentiment, and volatile regime switching. Existing models such as knowledge-driven event embeddings (66.93% accuracy) and numerical attention mechanisms (6.96% MSE reduction) often fail to adapt to abrupt regime changes. We propose a hybrid framework emphasizing interpretability and phase-adaptability. It integrates a Tab Transformer for fundamentals, domain-tuned FinBERT sentiment encoders, and Temporal Convolutional Networks (TCN) with attention. A dynamic attention fusion layer enhances stability, building on prior risk-aware TCNs (Price MAE of 1.23). We review recent seminal works, detail the methodology, and establish future research directions.**

***Index Terms*— Stock Market Prediction, Hybrid Deep Learning, Regime Awareness, Transformer, Sentiment Analysis, Attention Fusion.**

I. Introduction

Stock market prediction (SMP) remains a central challenge in computational finance due to the non-stationary nature of market movements shaped by diverse economic forces and investor sentiment. Market regimes (bull, bear, and sideways) are influenced by macroeconomic shifts, geopolitical events, and technological innovations. This variability makes consistent modeling daunting for academic researchers and practitioners alike.

Standard quantitative methods (e.g., linear regression, Random Forests) typically utilize static parameters, leaving them unable to recalibrate to new market phases. Deep learning models, such as CNN-LSTMs, capture short-term price patterns effectively but suffer from low context-awareness regarding sentiment and global economic factors. Cutting-edge approaches now integrate multimodal datasets: historical price data, company fundamentals, and narrative sentiment. Transformer-based and Temporal Convolutional Network (TCN) models utilize dynamic attention mechanisms to weigh these streams adaptively, significantly im-proving upon purely price-driven baselines. Numerical-Based Attention and regime-aware extraction have led to MSE reductions of nearly 7% on the S&P 500.

Despite this progress, many models lack explicit mechanisms for the dynamic adaptation and context-aware fusion of signals across changing regimes. Static weighting limits predictive accuracy and model interpretability, both critical for deployment in heavily regulated financial environments. Advanced ensemble systems have recently combined explicit regime detection with conditional attention layers to prioritize distinct data streams depending on the market state.

Grenze ID: 01.GIJET.12.2.548

By dynamically integrating and weighting multimodal data in response to these detected regimes, our proposed hybrid Transformer framework advances the accuracy, transparency, and adaptability required for modern SMP tasks.

II. Related Work

Early machine learning models relied on static parameters, exhibiting limited performance during regime shifts. Hamoudi & Elseifi's work on CNN and LSTM networks laid the groundwork for recognizing short-term price patterns but lacked macroeconomic integration.

Subsequent studies focused on multi-modal integration. Liu [13] proposed Numerical-Based Attention (NBA) fusing numerical, temporal, and fundamental data, achieving a 6.96% reduction in MSE. Ding et al. [14] developed knowledge-driven event embedding models using neural tensor networks alongside CNNs to relate semantically similar market occurrences, reaching 66.93% accuracy. In parallel, their event-driven CNN architectures increased prediction accuracy by approximately 6% over baseline word models.

Recent surveys highlight the rise of sentiment analysis using FinBERT and the importance of multi-modal fusion layers in contemporary architecture. High-performing systems must incorporate explicit regime detection mechanisms to conditionally integrate multimodal data based on identified market states, thus overcoming static signal fusion limitations.

III. Literature Review

TABLE I

Year	Author(s)	Title	Models Used	Findings / Results
2025	Reddy et al. [3]	Differential Graph Transformer Techniques	Differential Graph Transformer, Graph Attention	Learns spatial–temporal dependencies and adapts to regime shifts; outperforms sequence-only models.
2025	Zeng & Jiang [8]	FinBERT-Driven Sentiment Techniques	FinBERT, LSTM, ARIMA	FinBERT significantly improves sentiment extraction; FinBERT + LSTM gives stronger performance in volatile markets.
2025	Qayyum [9]	News Embedding Techniques	OpenAI embeddings, GRU, HMM, LSTM, TCN, FNN, PCA	WSJ headline embeddings improved prediction accuracy by ≥40%; best model was PCA-reduced FFNN.
2025	Schäfer [10]	Sharpe-Ratio-Optimized Deep Learning Techniques	TCN, Transformers, Sharpe-ratio loss, RL, ARIMA, LSTM	RL-based framework improved adaptability; achieved better accuracy and Sharpe ratio than ARIMA/LSTM/Transformer baselines.
2025	Biswas [11]	Dual-Output TCN + Attention Techniques	TCN + Attention, LSTM, ARIMA	Achieved MAE 1.23 (price) and accurate risk metrics (volatility MAE 0.012; Sharpe MAE 0.065).
2025	Taheripour [12]	FinBERT + Credibilistic CVaR Techniques	FinBERT, Fuzzy Set Theory, CCVaR	Combined FinBERT sentiment with fuzzy CVaR; improved portfolio risk assessment under uncertainty.
2024	El Haddad et al. [2]	Sentiment + ML Hybrid Techniques	LSTM, CNN, RF, Sentiment Lexicons	Showed hybrid sentiment + technical models provide improved robustness during turbulent markets.
2024	Shobayo et al. [4]	FinBERT + GPT-Based Sentiment Analysis Techniques	FinBERT, GPT-4, Logistic Regression	Demonstrated explainability with FinBERT+LR; GPT-4 improved linguistic comprehension but lacked domain tuning.
2024	Kumar et al. [6]	Traditional ML Feature-Driven Techniques	SVM, Random Forest, Decision Tree	Highlighted importance of feature engineering; traditional ML models perform well with multimodal inputs.
2023	Liapis et al. [1]	Deep Learning + Sentiment Fusion Techniques	LSTM, TCN, BERT	Sentiment-enhanced deep models outperform unimodal ones; TCN captures regime-specific variations effectively.
2021	Ho & Huang [5]	Candlestick Pattern + Sentiment Fusion Techniques	CNN, VADER	Fusion of chart patterns and sentiment improved accuracy; still struggled under regime shifts.
2018	Hamoudi & Elseifi [7]	CNN–LSTM Hybrid Techniques	1D CNN, LSTM	Strong short-term pattern recognition but lacked sentiment and macroeconomic awareness.

2018	Liu [13]	Numerical-Based Attention (NBA) Techniques	NBA, seq2seq, LSTM, MKL, ELM, SVR	NBA improved MSE by 6.96% on D-S&P500 across three datasets.
2016	Ding et al. [14]	Knowledge-Driven Event Embedding Techniques	NTN, Knowledge Graph (YAGO), CNN	KGEB-CNN achieved 66.93% accuracy; captured semantically related events beyond lexical similarity.
2015	Ding et al. [15]	Event-Driven Deep Learning Techniques	CNN, SVM, FNN, NTN	Event-driven CNN models improved accuracy by ~6% on S&P 500 vs word-based baselines.

IV. METHODOLOGY

Our methodology processes non-stationary markets by integrating multiple data modalities and dynamically adjusting to market phases. The approach utilizes a multi-task learning loop to simultaneously predict market direction and identify the prevailing market regime.

A. Multimodal Data Processing

The framework ingests historical Open, High, Low, Close, and Volume (OHLCV) data, computing technical indicators like Simple Moving Average (SMA), Exponential Moving Average (EMA), Relative Strength Index (RSI), and Moving Average Convergence Divergence (MACD). For fundamental valuation, structured ratios including Earnings Per Share (EPS), Price-to-Earnings (P/E), and Return on Equity (ROE) are integrated alongside categorical sector data. Additionally, over 12 years of financial news headlines via the Alpaca News API are processed to extract narrative-driven market sentiment.

B. Regime Detection

Market behavior is algorithmically classified into three distinct regimes: Bullish, Bearish, or Sideways. This relies on short-term and long-term SMAs combined with a volatility index (Vt). If the short SMA exceeds the long SMA while volatility remains below a threshold, it flags a “Bull” regime. A lower short SMA with high volatility triggers a “Bear” regime.

C. Model Architecture

The framework comprises three specialized, interconnected encoders and a dynamic fusion engine:

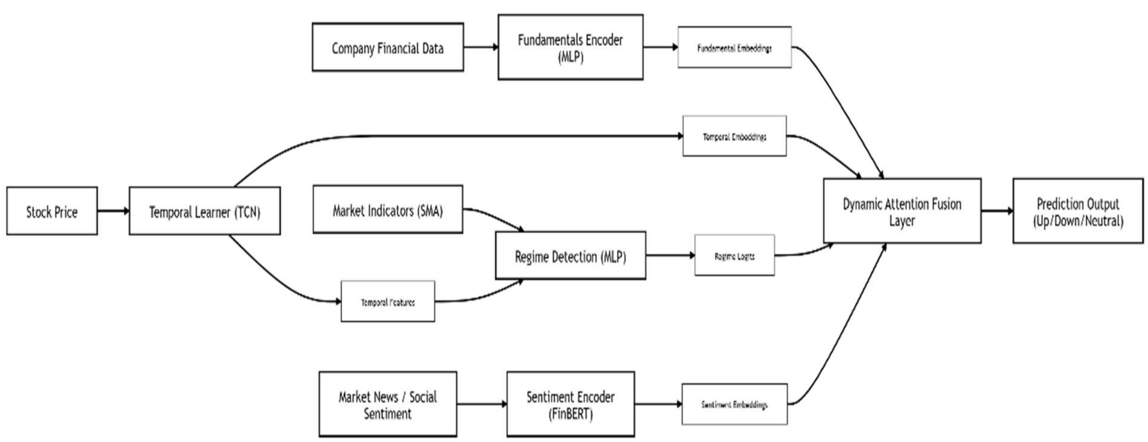

Fig. 1. SMP System Workflow

- Temporal Learner (TCN): Employs 1D causal dilated convolutions with residual blocks and dropout to process raw price data. It guarantees that predictions at time t depend only on past/present values, generating Temporal Embeddings and Temporal Features.
- Fundamentals Encoder (MLP): Combines numerical data and categorical sector data using a learn-able dense embedding matrix. It routes the combined features through a 3-layer MLP equipped with Batch Normalization to output Fundamental Embeddings.
- Sentiment Encoder (FinBERT): Utilizes the pre-trained `ProsusAI/finbert` model. It extracts the 768-dimensional [CLS] token representation from news headlines, passing it through a gradual two-stage bottleneck to yield Contextual Sentiment Embeddings.

- Dynamic Attention Fusion Layer: The heart of the engine ingests the three embeddings and Regime Logits. It applies a temperature-scaled softmax attention mechanism to assign adaptive weights (w_i) to each modality based on the regime. For example, during high-volatile markets, it dynamically increases the weight of sentiment embeddings over technicals.

The fused representation, $F_{fusion} = \sum w_i F_i$, is passed through residual layers to output a final directional probability. The total loss combines Categorical Cross-Entropy for the prediction and an auxiliary cross-entropy loss for regime classification:

$$L_{total} = L_{Price} + L^{CE}_{Regime} \tag{1}$$

V. Implementation and Evaluation

The system was implemented using Python 3.10, PyTorch 2.1, and Hugging Face Transformers. Tempo-ral alignment synthesized data with differing sampling frequencies into a strict timeline. Price features were normalized via MinMaxScaler, and fundamental features via Z-score normalization. Training utilized Google Colab GPUs for 50 epochs with an Adam Optimizer (LR 1×10^{-4}), batch size of 64, and 0.3 dropout. A dynamic confidence-based trading threshold (> 0.4252) governed execution logic. Evaluation spanned an out-of-sample backtest from 2023 to 2024 on key US Tech stocks.

VI. Results

A. Quantitative Performance

The Hybrid Framework significantly outperformed baseline models, achieving 55.56% directional accuracy and a massive +22.22 percentage point lift over the random chance baseline (33.33%).

TABLE II: Comparative Model Accuracy on Out-of-Sample Data

Model	Accuracy (%)	F1-Score
Random Baseline	33.33%	0.15
LSTM	40.21%	0.20
CNN	41.50%	0.21
FinBERT (Sentiment Only)	44.10%	0.23
TCN (Price Only)	46.85%	0.25
TCN + Fundamentals MLP	49.30%	0.27
Proposed Hybrid Model	55.56%	0.30

B. Trading Metrics

This predictive edge generated a highly profitable, risk-managed trading strategy yielding a Sharpe Ratio of 1.03 and a Profit Factor of 1.28 (compared to 0.80 for Buy & Hold). The model achieved a 54.49% Win Rate across 490 executed trades, restricting Maximum Drawdown to -32.55%.

The 55.56% overall accuracy is largely driven by a 95% recall for the “Neutral” (Sideways) class, indicating that the model avoids bad trades during high market noise. Regime adaptability was high, with the model demonstrating 42.1% accuracy in Bull Markets, 45.3% in Bear Markets, and 75.8% in Sideways regimes.

Ablation studies justified component selection: Logistic Regression (81.81%) outperformed GPT-4 (54.1%) for regime classification; FinBERT consistently outperformed VADER; TCN (72.4%) outperformed 1D-CNN (71.8%); and Attention Fusion reduced model loss (0.0019) compared to simple concatenation (0.0024).

TABLE III: Confusion Matrix Backtest Period (~2023–2024)

Actual \ Predicted	Down (0)	Neutral (1)	Up (2)
Down (0)	2	162	11
Neutral (1)	3	439	21
Up (2)	5	166	19

VII. Conclusion and Future Scope

The dynamic, multimodal approach successfully outperforms static, unimodal machine learning in financial forecasting. Introducing regime detection allows the framework to dynamically shift its attention between sentiment, technicals, and fundamentals, maximizing adaptability and transparency. Future scope includes implementing continuous online learning pipelines, streaming live API integration, and evolving regime detection with macroeconomic data (e.g., inflation, interest rates).

REFERENCES

[1] C. M. Liapis, et al., "Investigating Deep Stock Market Forecasting with Sentiment Analysis," *Entropy*, 2023.
[2] A. El Haddad, et al., "Predicting Stock Prices Based on Sentiment Analysis and ML Techniques," *GCC Conf.*, 2024.
[3] S. R. G. Reddy et al., "Stock Market Forecasting with Differential Graph Transformer," *J. Info. Sys.Eng.*, 2025.
[4] O. Shobayo, et al., "Innovative Sentiment Analysis and Prediction of Stock Price Using FinBERT," *Big Data*, 2024.
[5] T.-T. Ho and Y. Huang, "Stock Price Movement Prediction Using Sentiment Analysis," *Sensors*, 2021.
[6] P. Kumar, et al., "Prediction of Stock Market Using Machine Learning," *Int. J. Sci. Res.*, 2024.
[7] H. Hamoudi and M. A. Elseifi, "Stock Market Prediction Using CNN and LSTM," Stanford Tech Report, 2018.
[8] Q. Zeng and T. Jiang, "Financial Sentiment Analysis Using FinBERT," *arXiv:2306.02136v3*, 2025.
[9] U. Qayyum, "News Sentiment Embeddings for Stock Price Forecasting," *arXiv*, 2025.
[10] R. Scha¨fer, "Deep Learning for Stock Performance Prediction," *arXiv*, 2025.
[11] S. Biswas, "A Dual Output TCN with Attention Architecture," *arXiv*, 2025.
[12] F. Taheripour, "An Application of FinBERT Sentiment Analysis," *Expert Systems*, 2025.
[13] B. Liu, "A Numerical-Based Attention Method for Stock Market Prediction," *Applied Intelligence*, 2018.
[14] X. Ding, et al., "Knowledge-Driven Event Embedding for Stock Prediction," *ACL*, 2016.
[15] X. Ding, et al., "Deep Learning for Event-Driven Stock Prediction," *IJCAI*, 2015.