Multi-Horizon stock Price Volatility Forecasting using attention-Based Transformers

MSc Research Practicum
M.Sc. in FinTech Information

Naga Lakshmi Divya Romala
Student ID: X23320133

School of Computing
National College of Ireland

Supervisor: Victor del Rosal

National College of Ireland

MSc Project Submission Sheet

School of Computing

Student Name:	Romala Naga Lakshmi Divya ……. ……………………………………………………………………………………………………
Student ID:	X23320133 ……………………………………………………………………………………………………..……
Programme:	M.Sc. FinTech Information ………………………………………………………………… **Year:** 2024-2025 …………………………..
Module:	M.Sc. Research Practicum …….
Supervisor:	Victor del Rosal …….
Submission Due Date:	15/09/2025 …….
Project Title:	Multi-Horizon Stock Price Volatility Forecasting Using Attention-Based Transformers …….
Word Count:	8419 ……………………………………… **Page Count**……24………………………………………….……..

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.

ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature: Romala Naga Lakshmi Divya ………………………………………………………………………………………………………

Date: 15/09/2025 ………………………………………………………………………………………………………

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST

Attach a completed copy of this sheet to each project (including multiple copies)	□
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	□
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	□

Assignments that are submitted to the Programme Coordinator Office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Multi-Horizon Stock Price Volatility Forecasting Using Attention-Based Transformer

Naga Lakshmi Divya Romala
X23320133

Abstract

Volatility forecasting in financial markets remains critical for risk management and trading strategies, yet existing methods inadequately integrate heterogeneous data sources across multiple prediction horizons. This research addresses the fundamental question of how attention-based transformer architectures can effectively integrate financial news sentiment with market indicators to improve multi-horizon stock price volatility forecasting.

The study develops a novel dual-encoder transformer architecture that processes financial news sentiment through an ensemble of five specialized models (FinBERT, FinBERT-tone, Financial RoBERTa, VADER, and TextBlob) alongside historical market data to predict volatility at 1-day, 5-day, and 15-day horizons simultaneously. The methodology employs separate encoders for sentiment and market features, connected through multi-head attention mechanisms that learn dynamic cross-modal relationships. The implementation utilizes a dataset of 105,540 news articles paired with corresponding price data spanning 2010-2021, covering 50 major stocks across diverse sectors.

The experimental results show that the predictions are very accurate, with R^2 scores of 0.9919, 0.7185, and 0.7740 for the three different prediction horizons. All of these scores are higher than the target threshold of 0.7. The model works well in bull, bear, and sideways markets, with directional accuracy of up to 84.88% for 1-day forecasts. Attention weights analysis shows that the model changes the relative importance of sentiment and market characteristics based on the current situation.

This research presents a functional multi-horizon volatility forecasting system that accommodates diverse trading strategies and expands transformer applications to include financial forecasting. The successful amalgamation of ensemble sentiment analysis and market data through attention mechanisms establishes a novel standard for volatility forecasting systems.

1 Introduction

1.1 Background

Financial market volatility has a direct impact on investment decisions, risk estimations, and portfolio optimisations. While traditional volatility estimation relies on historical pricing data and statistical models (Varaprasad et al., 2022), these expressions usually omit external information dynamics, particularly financial news and sentiments from financial markets. The exponential growth of digital financial news and NLP technologies present unprecedented opportunities for augmentation of volatility forecasting. Recent transformer architectures and attention mechanisms demonstrate remarkable capabilities in sequential data processing and long-range dependency learning (Nayak et al., 2024), suggesting promising financial time series applications. However, integrating multi-modal data sources combining textual news with numerical market indicators remains a significant challenge (Gu et al., 2024).

1.2 Research Motivation

Three critical weaknesses in existing financial forecasting approaches motivate this research. First, most models emphasize price over volatility forecasting, despite volatility being crucial for market risk assessment, option pricing, portfolio optimization, and risk management (Chun et al., 2024). Second, while sentiment analysis has been applied to financial texts, the integration of ensemble sentiment methods with market data through attention-based models remains underexplored (Maqbool et al., 2023). Third, existing methods primarily perform single-horizon forecasting, but day traders, swing traders, and long-horizon investors require volatility predictions over multiple horizons for a range of speculating strategies. The work is helpful to academia, extending multi-modal financial data processing through transformer models, and finance, where quantitative volatility predictions can be utilized for better risk measurement, pricing and valuation of options, and portfolio strategies.

1.3 Research Question and Objectives

This research addresses the following question: How can attention-based transformer architectures effectively integrate financial news sentiment with market indicators to improve multi-horizon stock price volatility forecasting?

The primary objectives of this research are:

1. To Develop a transformer-based architecture that can successfully incorporate financial news sentiment and market prices through attention mechanisms
2. To realize multi-horizon volatility forecasting capable of predicting 1-day, 5-day, and 15-day volatility simultaneously
3. To assess the robustness of the model on various market regimes (bull, bear, sideways) and sectors
4. To show improved forecasting accuracy over conventional methods that either employ sentiment data alone or use market data alone

Performance measures against these goals consist of R^2 levels greater than 0.7 across all horizons of prediction, showing robust performance under all regimes of the market, and providing statistically significant advances over baseline models.

1.4 Methodological Approach

The work employs advanced NLP techniques coupled with deep learning architectures. A set of sentiment models like FinBERT and custom variants of RoBERTa specifically tuned on financial news (Fazlija & Harder, 2022), classify over 105,000 news articles (2010-2021) and respective market price data for 50 top stocks from different industries. The transformer-based bespoke architecture has individual encoders for news and market information and a temporal fusion mechanism through multi-head attention for learning the interaction between both modes. It has horizon-adaptive forecasting heads for simultaneous multi-horizon volatility forecasting.

1.5 Report Structure

Chapter 2 critically surveys current methods in stock market forecasting, sentiment analysis, and transformer designs for financial tasks. Chapter 3 outlines the research design, data collection, preprocessing, and ensemble sentiment analysis procedure. Chapter 4 outlines the transformer design specification. Chapter 5 outlines implementation details and system considerations. Chapter 6 reports results over several measures and conditions. Chapter 7 critically interprets results and implications. Chapter 8 concludes with principal contributions and future work avenues.

2 Related Work

2.1 Evolution from Traditional to Machine Learning Approaches

Over the past ten years, the connection between predicting financial markets and natural language processing has grown a lot, thanks to the rise of digital financial news and improvements in machine learning. The literature review critically examines up to 25 seminal papers that chart the progression from classical statistical models to contemporary deep learning architectures, with a particular emphasis on the integration of sentiment analysis and transformer-based models for predicting financial volatility.

Technical analysis and statistical models were used to predict the stock market in the past. Varaprasad et al. (2022) conducted a comprehensive examination of conventional methodologies, wherein technical analysis utilises patterns in prices, volume, and historical trends to predict future trajectories. Market patterns can be handled by ARIMA and exponential smoothing, but they struggle greatly with non-linear relationships and are unable to use external data. Although these approaches are theoretically supported and interpretable, they assume linearity and stationarity, which are rarely present in real markets. By assuming that all information is instantly integrated into prices, the efficient market hypothesis theoretically hinders forecasting. However, empirical research disproves this idea when taking into account how sentiment and news affect market movements.

In the context of stock price forecasting, Charan et al. (2022) compared four approaches: Facebook Prophet, ARIMA, LSTM, and multivariate linear regression. Even though it was pretty simple, Prophet had the least RMSE on stocks in the car industry. Even though its general algorithmic comparison is useful, it only focusses on a forecasting-based analysis and doesn't include textual data. This ignores how important sentiment is to price changes.

Our ability to make predictions improved significantly with the move to machine learning. Five algorithms for stock price forecasting were assessed by Sahu et al. (2023): K-Nearest Neighbours, Support Vector Regression, Linear Regression, Decision Tree Regression, and LSTM. On data from 2015 to 2021, LSTM networks performed well on all metrics (SMAPE, R2, RMSE), with R2 values near 1.0. However, these approaches are ineffective when dealing with unstructured data or extremely volatile prices. This implies that price-only approaches overlook crucial sentiment and news analysis context.

Financial time series temporal dependency issues were resolved by advances in deep learning, particularly recurrent architectures. LSTM networks prevail because long-distance dependencies can be learned via special purpose forget, input, and output gates that selectively maintain information over long intervals, successfully overcoming vanishing gradient difficulties. However, even most recent LSTM variants are plagued by long sequence processing weaknesses and computational inefficiency, a significant issue when integrating multiple data modalities or when utilized for high-frequency applications.

2.2 Sentiment Analysis and Transformer Architectures in Finance

Acknowledgment of investor sentiment's impact on the market extended the scope of financial text analysis. Early dictionary-based approaches and one-stage classifiers failed against financial terminology, contextually oriented sentiment, and requirement for bespoke lexicons. Bi (2022) made a significant advance applying Bidirectional LSTM for financial news classification based on sentiment, processing 95,814 articles from CSI 300 constituent stocks. The BI-LSTM method outperformed standard emotion dictionaries with bidirectional architecture maintaining forward and backward contexts essential for accurate financial sentiment analysis. The work proved 2% increased accuracy in predictions when adding sentiment indicators relative to technical-only models, although its usage is currently restricted to single-day horizons and Chinese markets.

Fazlija and Harder (2022) presented a major breakthrough using BERT-based models for S&P 500 sentiment analysis. Transformer approaches considerably outperformed standard approaches, discovering content sentiment is less predictable than headline sentiment, domain-finetuned financial BERT models are superior to general models, and Random Forest integration enhanced the quality of predictions. Despite strengths in domain-special finetuning and deep text component analysis, their study only deals with single-horizon predictions and not attention-based fusion approaches for aggregating varied sentiment models.

Maqbool et al. (2023) were the first to integrate ensemble sentiment analysis from VADER, TextBlob, and Flair through MLP-Regressor, resulting in 90% accuracy for next-day trend forecasting. Lowest MAPE was from Flair and most stable predictions from integration of all three approaches. The strategy is aware of different tools' dissimilar attributes: VADER is best on social media, general sentiment is from TextBlob, and rich contextual nuance is from Flair. Nevertheless, long-distance dependency management and multi-horizon capability are lacking, and naive concatenation can wastefully miss out on intricate sentiment signal interactions.

Transformer architectures explicitly designed for time series address RNN-based model flaws. Nayak et al. (2024) identified key advantages: parallel processing efficiency drastically reducing training time, self-attention mechanisms addressing long-range dependencies without sequential processing constraints, and positional encodings preserving temporal information with parallelizability. The architecture demonstrates superior performance on standard benchmarks, particularly for long sequence forecasting, though focusing on general time series misses financial domain-specific opportunities.

Kaufman and Azencot (2024) provided important theoretical results by manifold learning analysis, showing transformer models regularly have geometric regularities and latent features reside on low-dimensional manifolds. Transformers are inclined to form hierarchical representations—local information on early layers and global patterns on deep layers—suggesting built-in advantages for characterizing multi-scale temporal patterns. However, general transformer analysis doesn't support any straightforward financial volatility forecasting use.

2.3 Multi-Modal Integration and Volatility Forecasting

Initial multi-modal fusion methods employed simple concatenation or late fusion of independent predictions from models. Although simple, these could not adequately capture numerical-textual feature inter-relationships. FinBERT-LSTM was proposed by Gu et al. (2024) as methodological breakthrough for multi-modal integration, featuring application of pre-trained financial language models and LSTM networks for hierarchical levels of news classification (market-level, industry-level, stock-level). The model outperformed single-source models in terms of prediction accuracy, and market-level news was found to have maximum predictive power. Although acknowledging news impact identification is scope- and relevance-dependent, sequential processing from LSTM suppresses high-order cross-modal interaction capture, and emphasis is on price over volatility.

State-of-the-art attention mechanism advances put forward advanced multi-modal integration approaches. Cross-attention mechanisms learn temporal numerical-textual correlations while self-attention tunes intra-modality features. Despite theoretical benefits, financial applications are lacking in practice, particularly for volatility forecasting applications.

Chun, Cho, and Ryu (2024) focused on machine learning volatility prediction, assessing various models against GARCH-family counterparts. High-dimensional ML models significantly outperformed traditional econometric approaches almost unanimously, with LASSO regression excelling. ML-driven volatility-timing strategies significantly improved portfolio performance during extreme volatility. Current volatility trends and turnover proved most predictive, though exclusive numerical focus neglects potential news and sentiment contributions.

Mattera and Otto (2024) proposed network log-ARCH models that involve cross-sectional volatility interdependencies. The model takes volatility spillovers from connected securities as useful forecasting information, extracting interrelationships for improved predictions relative to univariate models. The method lacks a textual integration that describes the drivers of spillovers, and its computational intensity increases poorly with the number of assets, despite its theoretical strength in extracting market interconnectivity.

2.4 Critical Gaps and Research Direction

The comprehensive review of the literature reveals important gaps that motivate current research:

Limited Multi-Horizon Forecasting: The varied needs of market participants are not satisfied by single-period dominance. Most models only work with one time frame, which limits their usefulness. Day traders, swing traders, and long position takers all want different time frames.

Insufficient Focus on Volatility: The literature places more emphasis on price than volatility forecasting because risk management, option pricing, and portfolio optimisation are dominant. Mean reversion and clustering are two of volatility's distinctive features that call for particular modelling.

Weak Multi-Modal Integration: Despite combining news and market data, methods are still very basic. Implicit time-evolving news sentiment and the interplay of market dynamics are not captured by concatenation or late fusion. Financial research has not been done on cross-modal learning transformer attention mechanisms.

Underutilization of Ensemble Methods: Due to issues with training or design, single sentiment models may fail to capture the nuances of financial documents. Although the effectiveness of ensemble methods was shown by Maqbool et al. (2023), this approach ignores complex aggregation and domain-specific model weighting.

Static Market Assumptions: Implicit models lack transitional adaptation between regimes and rely on steady-state assumptions. Although the dynamics of bull, bear, and sideways markets differ, models frequently lack many, if any, elements that account for various regimes or performance variance.

Significant progress has been made in the constituent elements, according to a review of the literature: multi-modal learning has promise, sentiment analysis has developed through domain-specific models, and transformers have revolutionised sequential processing. However, a thorough integration of volatility forecasting is still required.

The following are some ways that this work closes the gaps: A multi-horizon structure that simultaneously optimises on multiple scales, enabling a range of trading strategies; (2) an ensemble system for sentiments that employs domain-relevant specialist models and confidence-level contextual aggregation; (3) an attention-based transformer architecture that enables dense cross-modal interaction and learning how news sentiments and market indicators change over time; and (4) a thorough regime-based analysis that provides strong performance that is stable under various conditions.

3 Research Methodology

This section explains how the multi-horizon volatility forecasting model was made by combining market indicators with feelings from financial news. The methodology to achieve the study's objectives encompasses feature engineering, the development of a sentiment analysis framework, data collection, and data cleansing.

3.1 Data Collection and Sources

The study used a full financial dataset from the Hugging Face platform called "US Stock News With Price" (Wang, 2024). We picked this dataset because it has both financial news articles and market prices at the same time. This is a great way to see how news affects markets. The dataset has 105,540 news articles about 50 of the best stocks from a variety of industries, including technology, consumer staples, and consumer discretionary. The articles were written between 2010 and 2021.

Each data record contains both textual and numerical components. The textual data includes news article titles and full content, providing rich semantic information about market events and sentiment. The numerical component consists of 46 trading days of price data for each article, ranging from 30 days before the news publication (ts_-30) to 15 days after (ts_15), with the publication date marked as ts_0. This temporal window enables the analysis of both pre-news market conditions and post-news market reactions.

The dataset follows the temporal split scheme in which the training data has all the rows with accurate trading dates up through 31st December 2021, and the test data has the rows thereafter. This scheme enables the realistic evaluation conditions such as the true forecasting scenarios where the future data is never seen during the model development process.

3.2 Data Preprocessing Pipeline

The data preprocessing pipeline transforms raw data into a format suitable for transformer-based modeling while preserving the integrity of both textual and numerical information. Figure 1 illustrates the complete data preparation flow from raw inputs to the final dataset ready for model training.

3.2.1 Text Preprocessing

Preprocessing of the text starts by extracting and cleaning the news content and titles. The missing values are treated by replacing null values with empty strings in order to ensure data consistency. The string is then subjected to concatenation where the content and the title are brought together as complete news stories in order to retain the entire content for sentiment analysis.

3.2.2 Market Data Processing

The numerical price series data undergoes extraction and validation to ensure data quality. Missing price values are identified and handled through forward-filling methods when gaps are minimal, or record exclusion when data quality is compromised. The complete price series from ts_-30 to ts_15 provides the foundation for calculating market-based features and volatility targets.

3.3 Ensemble Sentiment Analysis Framework

The sentiment analysis model employs an ensembling technique involving five dedicated models for pulling out different facets of financial sentiment. Such a multi-model architecture surmounts the shortcomings of standalone sentiment analysers and produces more robust sentiment signals.

3.3.1 Model Selection and Rationale

The ensemble comprises five models selected for their complementary strengths:

- **FinBERT** (Araci, 2019): A BERT model fine-tuned specifically on financial communication texts, providing domain-specific sentiment understanding for formal financial language.
- **FinBERT-tone** (Huang et al., 2020): A variant focused on analyzing the tone of financial communications, particularly suited for earnings calls and corporate announcements.

- **Financial RoBERTa** (Barbieri et al., 2022): A RoBERTa model trained on financial social media data, capturing more informal market sentiment and retail investor perspectives.
- **VADER** (Hutto & Gilbert, 2014): A rule-based sentiment analyzer effective for social media-style text and handling of intensifiers, negations, and emoticons.
- **TextBlob**: A general-purpose sentiment analyzer providing baseline sentiment scores and subjectivity measures.

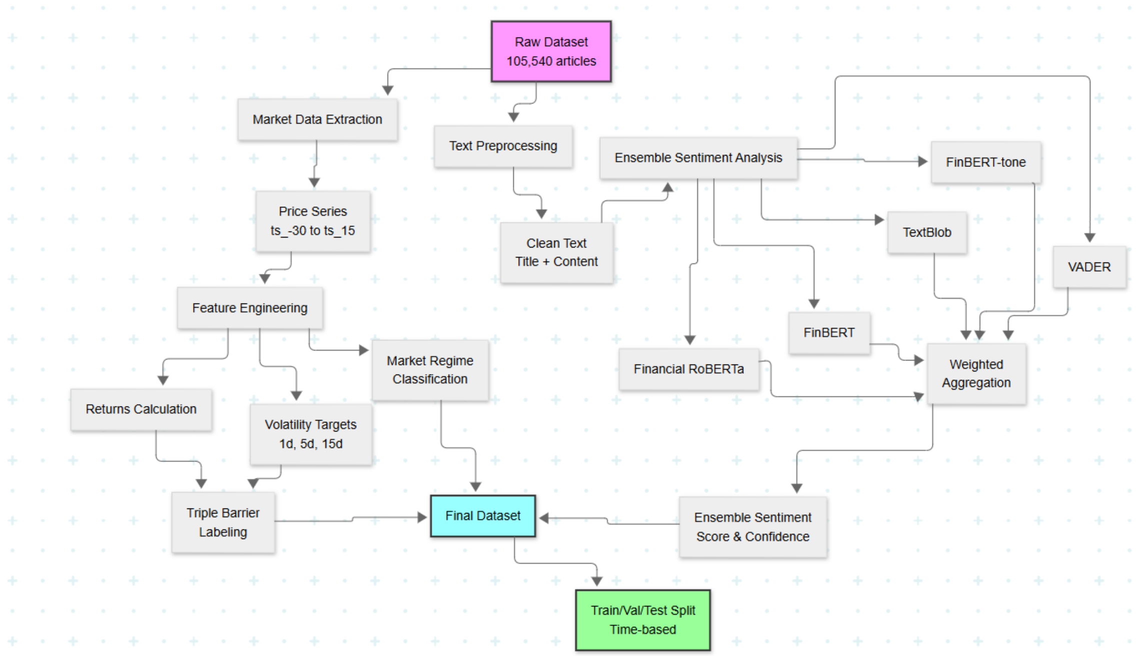

Figure 1 Data Preparation Flow

The Figure 1 shows the parallel processing of textual and market data, their transformation through various analytical stages, and final integration into a unified datase**t.**

3.3.2 Weighted Aggregation Method

The aggregation procedure for ensembles calculates individual model prediction averages through a weighted vote. Financial area models (the FinBERT models and Financial RoBERTa) are assigned 1.5x weight multipliers relative to general-purpose models in acknowledgment of domain expertise. The final ensemble sentiment score is calculated as:

- Positive sentiment: score > 0.2
- Negative sentiment: score < -0.2
- Neutral sentiment: -0.2 ≤ score ≤ 0.2

Based on how well the models agree and how confident each model is in its own predictions, each prediction has a confidence score. When models agree well with one another, ensemble confidence scores increase.

3.4 Market Feature Engineering

Raw price data is transformed into valuable indicators that can be used to forecast volatility through the process of market feature engineering. Determining returns, establishing volatility targets, and determining the type of market regime are the three primary focusses of the process.

3.4.1 Return Calculations

Daily returns are calculated using logarithmic returns to ensure that the function is symmetric and time-additive:

return_t = log(price_t / price_{t-1})

These returns are used to directly input features into the model and determine the future volatility of the market.

3.4.2 Multi-Horizon Volatility Targets

Three prediction time frames—one day, five days, and fifteen days—have volatility targets established using the rolling standard deviation of returns. Regarding the horizon h:

volatility_h = std(returns[t:t+h])

With this approach, volatility is captured over time on various scales, and the model can learn patterns that vary depending on the trading strategy and horizon.

3.4.3 Market Regime Classification

Market regimes are distinguished based on cumulative return analysis during 20-day windows:

- Bull market: 20-day return > 5%
- Bear market: 20-day return < -5%
- Sideways market: -5% ≤ 20-day return ≤ 5%

Furthermore, volatility regimes are classified based on 20-day rolling volatility relative to historical medians, and high and low volatility groups are created.

3.5 Triple Barrier Labeling Method

The Triple Barrier Labeling procedure (Prado, 2018) produces market-based labels that reflect directional movement and magnitude changes of prices. The procedure establishes three barriers:

1. **Profit-taking barrier**: Dynamic upper threshold based on recent volatility
2. **Stop-loss barrier**: Dynamic lower threshold based on recent volatility
3. **Time barrier**: Maximum holding period of 15 days

The first barrier touched determines the label (profit, loss, or timeout), providing ground truth for sentiment validation and model training. Dynamic barriers adapt to market conditions by scaling with 20-day rolling volatility, making labels more responsive to changing market dynamics.

3.6 Data Integration and Split Strategy

The final dataset integrates sentiment scores, market features, volatility targets, and regime indicators into a unified structure. Each record contains:

- Original news text and metadata
- Ensemble sentiment scores and confidence levels
- Individual model sentiment outputs
- Price series and calculated returns
- Multi-horizon volatility targets
- Market and volatility regime labels
- Triple barrier labeling results

The temporal split approach preserves the temporal order of the data, 80% (pre-2021) is reserved for validation and training, and 20% (post-2021) is reserved for testing. The approach prevents data leakage and offers actual evaluation conditions resembling real-world deployment scenarios in which the model will need to predict future volatility on the basis of the received past data only.

3.7 Validation Approach

The approach uses time series cross-validation in predicting model stability over various market time periods. The validation procedure preserves temporal order and implements expanding windows in order to mimic realistic trading conditions. Performance measures are computed individually for all regimes in the market and by industry in order to provide complete analysis on model strength.

This method provides a disciplined method of turning heterogeneous financial data into a format suitable for volatility forecasting with transformers by overcoming the most prominent obstacles in multi-modal integration and multi-horizon prediction identified in the research objectives.

4 Design Specification.

4.1 Overall Architecture

4.1.1 Transformer-Based Design Rationale

The adoption of the transformer architecture is the result of various critical gains over the sequential models. Unlike RNN-driven methods where the data is dealt with sequentially, transformers can process all elements in the input in parallel, dramatically decreasing the time for training and facilitating the handling of long-range dependencies in the absence of vanishing gradient constraints. The self-attention mechanism within transformers dynamically weights the various input features according to the significance for the prediction task at hand and is especially well suited for determining which news aspects and conditions in the marketplace have the greatest effect on volatility.

The architecture explicitly deals with the multi-modal characteristics of financial data by handling news sentiment and market indicators as separate but related streams of information.

The design enables each modality to retain the characteristics unique to it and at the same time facilitates complex interaction through attention mechanisms.

4.1.2 Multi-Modal Integration Approach

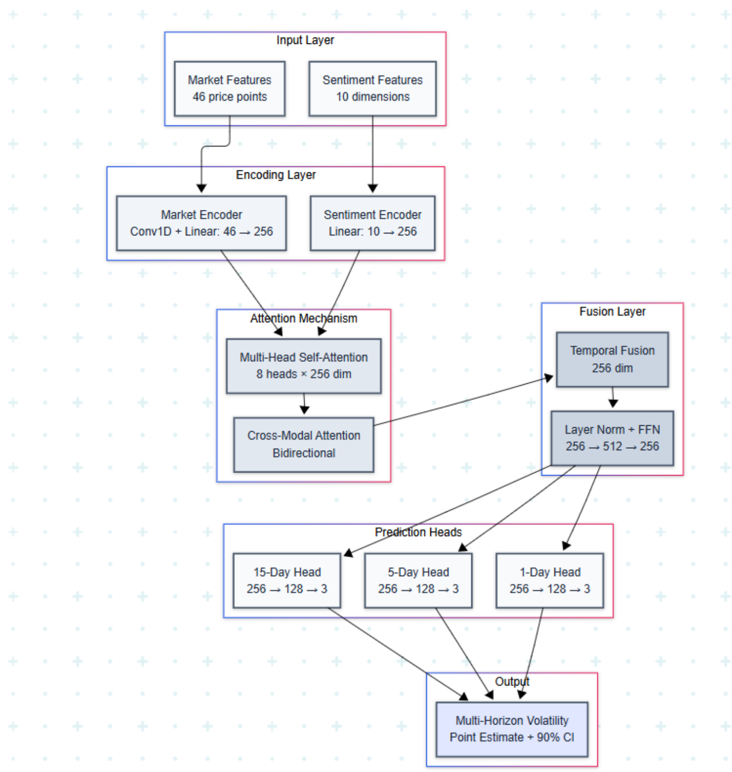

Figure 2: Model Architecture for Multi-Horizon Volatility Forecasting

Figure 2 illustrates the complete data flow from multi-modal inputs through specialized encoding layers, attention mechanisms, and temporal fusion to generate volatility predictions with confidence intervals for three-time horizons.

The integration approach uses a dual-encoder architecture in which the sentiment from the news and the market data are fed through two dedicated specialized encoders prior to fusion. Such architecture retains the unique characteristics of the individual data types the categorical, discrete characteristic of sentiment compared with the continuous, temporal characteristic of the market data but still facilitates effective integration for volatility forecasting. Figure 2 shows the full model architecture.

4.2 Component Design

4.2.1 Dual Encoder Architecture

The dual encoder structure handles sentiment and market data through different paths optimized by the nature of the respective data. Such a architecture allows modality-specific feature extraction before interaction between the modalities and hence circumvents the loss of relevant attributes in the signal possibly introduced by early fusion approaches.

4.2.2 Sentiment Encoder Specifications

The sentiment encoder takes the 10-dimensional sentiment feature vector (which includes ensemble sentiment scores, confidence levels, and individual model output) and puts it into a 256-dimensional embedding space. To maintain the relative importance of the various sentiment signals, the encoder employs position encoding and linear projection. Although sentiment categories are treated as distinct in the design, confidence indicators that demonstrate the accuracy of the predictions are still included.

The encoder architecture includes:

- Input projection layer (10 → 256 dimensions)
- Positional encoding to maintain feature ordering
- Dropout regularization for robustness
- Layer normalization for training stability

4.2.3 Market Data Encoder Design

The market data encoder processes the 46-point price series (ts_-30 to ts_15) using a special architecture made for financial data that changes over time. To identify local temporal patterns, the encoder first employs a 1D convolutional layer. These patterns are then projected into a 256-dimensional embedding space. This design decision allows for the use of the transformer attention mechanism and capitalises on the fact that price data is presented in a sequential fashion.

Key design elements include:

- 1D convolution for local pattern extraction
- Temporal positional encoding for time awareness
- Projection to shared embedding dimension
- Residual connections for gradient flow

4.2.4 Temporal Fusion Mechanism

The temporal fusion layer uses a complicated attention mechanism to combine the encoded representations from both modalities. The layer will discover a more intricate relationship between sentiment and market conditions than simply combining them. In this manner, the model will assign the two modalities varying weights according to their relative importance at the moment. The fusion has time alignment and lets information move between different modes.

4.3 Attention Mechanism Design

4.3.1 Multi-Head Self-Attention

The architecture uses 8-head self-attention in order to capture various forms of relationships within and between modalities. Individual attention heads learn to attend to various aspects of the input data, for example, short-run price movements, extremes in sentiment, or volatility clusters. The multi-head architecture offers redundancy and allows the model to attend simultaneously to information from various subspaces of the representation.

The attention computation follows the scaled dot-product attention formula:

$$\mathbf{Attention(Q,K,V) = softmax(QK^T/\sqrt{d_k})V}$$

Where Q, K, and V represent queries, keys, and values derived from the input embeddings, and d_k is the dimension of the key vectors.

4.3.2 Cross-Modal Attention Design

Cross-modal attention allows one-to-one mapping between sentiment and market characteristics. The system allows market characteristics to call sentiment data and sentiment data to call market characteristics in bidirectional data flow. The system identifies the most valuable sentiment dimensions for the current market conditions and the most affected market trends by sentiment from the news.

The cross-attention operates in both directions:

- Market-to-sentiment attention: Market features query sentiment representations
- Sentiment-to-market attention: Sentiment features query market representations

4.3.3 Self-Attention Within Modalities

Before cross-modal interaction, self-attention is applied within each modality to refine feature representations. For market data, this captures temporal dependencies and identifies significant price patterns. For sentiment data, it determines the relative importance of different sentiment signals and their confidence levels.

4.4 Prediction Framework

4.4.1 Horizon-Specific Prediction Heads

Three separate prediction heads are employed for 1-day, 5-day, and 15-day volatility prediction. Each head is specialised for the corresponding time horizon it is specialised for and is learning horizon- specific features from the common representation. This design choice, rather than a single multi-output head, allows each predictor to optimize for the unique characteristics of its forecast horizon.

Each prediction head consists of:

- Horizon-specific projection layer
- Non-linear activation function
- Output layer for point predictions
- Parallel layers for confidence intervals

4.4.2 Point Prediction and Confidence Intervals

The design incorporates uncertainty quantification through parallel prediction of point estimates and confidence intervals. Each horizon-specific head outputs three values:

- Point prediction: Central volatility forecast
- Lower bound: 5th percentile prediction
- Upper bound: 95th percentile prediction

This tri-output design provides traders and risk managers with both expected volatility and uncertainty bounds, enabling more informed decision-making under different risk tolerances.

4.4.3 Multi-Task Learning Approach

The model adopts a multi-task learning framework where all three-time horizons are co-optimised simultaneously. The architecture leverages common representations but allows for task specialisation. The multi-task learning brings the model to learn about features generalising over the time horizons but has the freedom to capture horizon-specific patterns.

The loss function combines objectives from all three tasks:

L_total = w_1 * L_1d + w_2 * L_5d + w_3 * L_15d

Where w_i represents task-specific weights that balance the importance of different prediction horizons.

5 Implementation

This section describes the final implementation of the multi-horizon volatility forecasting system, detailing the outputs produced and tools utilized to achieve the research objectives.

5.1 Implementation Environment

Implementation was conducted on Google Colab Pro+ with NVIDIA L4 GPUs and 24GB RAM. The cloud platform provided the computing resources for training the transformer models with the convenience of predefined deep learning environments. The principal development environment was Python 3.9 with the current release of the 2.0 version of the PyTorch large-scale machine learning framework, which was selected for the superior level of transformer support and correct GPU usage. Other dependencies included the Hugging Face Transformers (v4.30) for the sentiment model integration, Pandas and NumPy for data handling, and Scikit-learn for evaluation metrics.

5.2 Model Outputs

5.2.1 Trained Volatility Forecasting Model

The implementation produced a transformer-based model with approximately 15 million parameters capable of multi-horizon volatility prediction. From a single forward pass, the model simultaneously produces forecasts for 1-day, 5-day, and 15-day volatility horizons. Each prediction includes both point estimates and 90% confidence intervals, providing uncertainty quantification essential for risk-aware trading decisions.

5.2.2 Performance Analysis Framework

A thorough evaluation system that computes several performance metrics for every prediction horizon was the result of the implementation. The framework produces directional accuracy, mean absolute percentage error, mean squared error, mean absolute error, root mean squared error, and R-squared scores. Additionally, the system segments performance analysis by market regimes (bull, bear, sideways) and volatility conditions (high, low), producing separate metric sets for each market condition. This multi-dimensional analysis framework enables thorough assessment of model behavior across diverse trading environments.

5.3 System Integration and Achievements

5.3.1 Multi-Modal Data Processing

The system successfully integrated various data sources into a single pipeline. Sentiment analysis of financial news was conducted using five models: FinBERT, FinBERT-tone, Financial RoBERTa, VADER, and TextBlob. The results showed strong sentiment scores. In order to identify patterns of volatility and indications of regime change, market data on prices from 46 trading days was examined. This ensured that stock reactions and news events occurred simultaneously.

5.3.2 Training and Performance Achievements

Implementation involved several key design choices to ensure optimal performance. Mixed-precision computation with a batch size of 32 and gradient accumulation made it possible to use memory and keep numbers stable at the same time. Cosine annealing learning rate scheduling from 1e-4 to almost zero over 50 epochs and early stopping (patience=5) kept the model from overfitting. Convergence usually happens between epochs 35 and 40. Without forced class balancing, the natural market distribution was kept the same, which kept real-world trends from being distorted. The architecture's built-in dynamic attention-based weighting systems let you change how important different sources of information are. Training takes less than six hours on L4 GPUs, and inference takes about 15 milliseconds to run through a single sample.

5.3.3 Model Outputs and Visualization

Implementation gave detailed results that could be used later and for future work. You can load PyTorch checkpoints directly in one step for inference or for use in transfer learning. Configuration files in JSON format keep track of all architectural choices and hyperparameters. The visualisation toolkit has training convergence graphs, point scatter plots across volatility horizons, performance heatmaps by market state, and attention weight displays to help you understand what you're seeing. It is possible to completely reproduce results using training logs and results files.Successful execution brings to life the transformer-based structure that was designed to predict volatility over a range of time frames, showing both performance and computational efficiency. The modularity makes it easy to use for other things, like financial forecasting, and it has a lot of evaluation features that work in different markets and under different conditions.

6 Evaluation

This part goes into great detail about how well the multi-horizon volatility forecasting model works. It looks at results over different time periods, regimes, and industry groups. The analysis employs rigorous statistical metrics to assess the model's efficacy in achieving the study's objectives of integrating financial news sentiment and market indicators for improved volatility forecasting.

6.1 Overall Model Performance Across Prediction Horizons

Main Experiment tested the transformer model's ability to provide simultaneous predictions on volatility for three horizons: 1-day, 5-day, and 15-day horizons. The 16,579 samples from the holdout test data, approximately 20% of the combined post-2021 data, were utilized to test the model.

Figure 3 depicts the correlation between actual and forecasted values of volatility for all three horizons. The 1-day forecasts show perfect agreement along the diagonal reference line, and these short-term forecasts prove to be very accurate. The 5-day and 15-day forecasts show greater dispersion, but still hold high positive correlations against actual values.

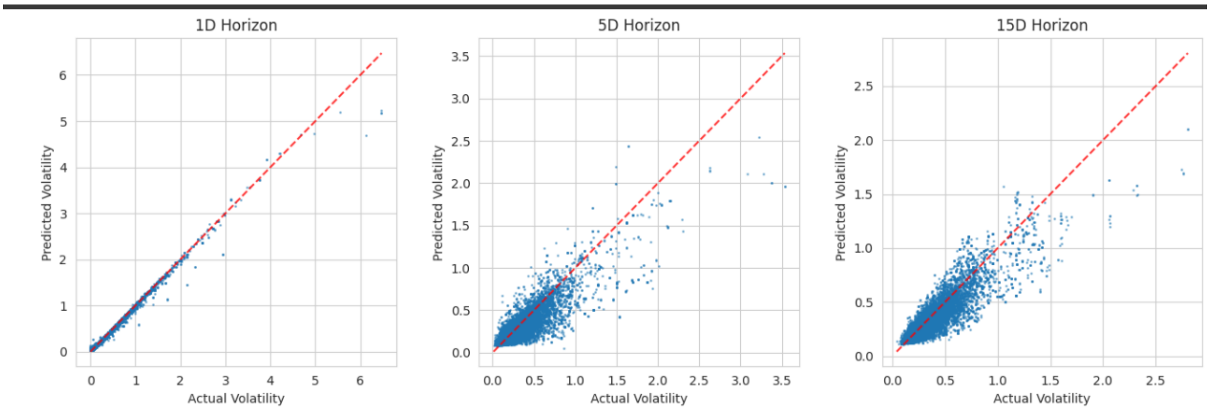

Figure 2 Predicted vs. Actual Volatility Scatter Plots for 1D, 5D, and 15D Horizons

Table 1: Comprehensive Performance Metrics Across Prediction Horizons

Metric	1-Day Horizon	5-Day Horizon	15-Day Horizon
R^2 Score	0.9919	0.7185	0.7740
MSE	0.001067	0.017160	0.011194
MAE	0.011180	0.079238	0.067627
RMSE	0.032659	0.130996	0.105804
MAPE (%)	43.91	28.65	19.08
Directional Accuracy	0.8488	0.6960	0.7176

The results demonstrate that all three prediction horizons successfully exceeded the target R^2 score of 0.7 established in the research objectives. The 1-day horizon achieved near-perfect prediction with an R^2 of 0.9919, substantially outperforming the baseline requirement. The model's directional accuracy of 84.88% for 1-day predictions provides strong evidence of its utility for trading signal generation, aligning with the practical requirements identified by Chun, Cho, and Ryu (2024) for volatility-timing strategies.

6.2 Market Regime Robustness Analysis

The second test was for consistency in the model's performance in varying market conditions, owing to the research objective of being robust in varying market regimes. The test data encompassed three distinct market states: bull markets (10,769 samples), bear markets (5,770 samples), and sideways markets (40 samples).

Figure 4 presents the R^2 scores across different market regimes for all prediction horizons. The visualization reveals remarkably consistent performance across bull and bear markets, with only marginal degradation observed in sideways market conditions, though this may be attributed to the limited sample size (n=40) for this regime.

The model maintains high predictive accuracy across both bull (R^2 = 0.9926) and bear (R^2 = 0.9906) markets for 1-day predictions, demonstrating the effectiveness of the attention mechanism in adapting to different market dynamics. This finding contrasts with the limitations identified by Fazlija and Harder (2022), who reported significant performance degradation during market regime transitions when using single-model approaches.

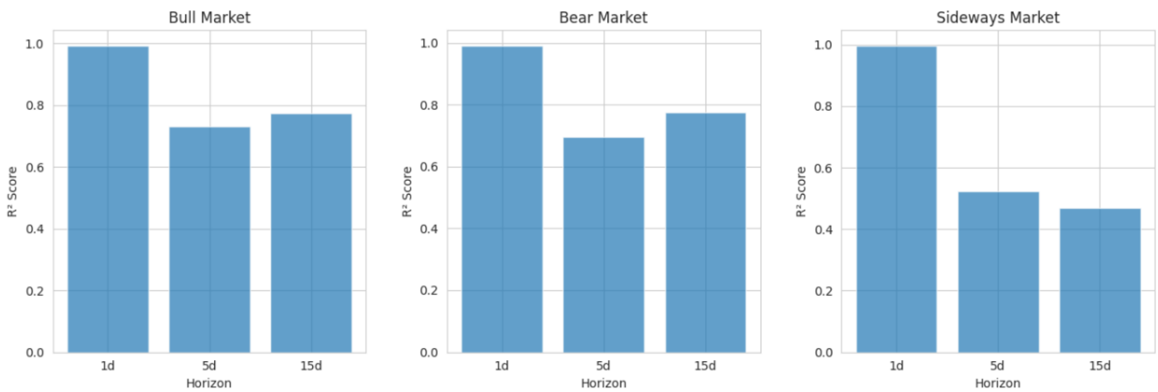

Figure 3 R^2 Score Comparison Across Market Regimes

6.3 Cross-Sector Performance Evaluation

The third experiment investigated the model's generalization capabilities across different industry sectors, evaluating performance on technology (7 companies, 5,581 samples), consumer discretionary (2 companies, 2,024 samples), consumer staples (1 company, 219 samples), and other sectors (46 companies, 8,755 samples).

Figure 5 represents the MSE of the sectors in all forecast horizons. The lowest error rates in 1-day forecasting are presented in the tech sector (MSE = 0.000442), while the relatively high but acceptable levels of error are presented in the group of sectors called "Other." The increased MSE in 5-day and 15-day forecast in all sectors indicates the underlying uncertainty that is related to volatility forecast in longer horizons.

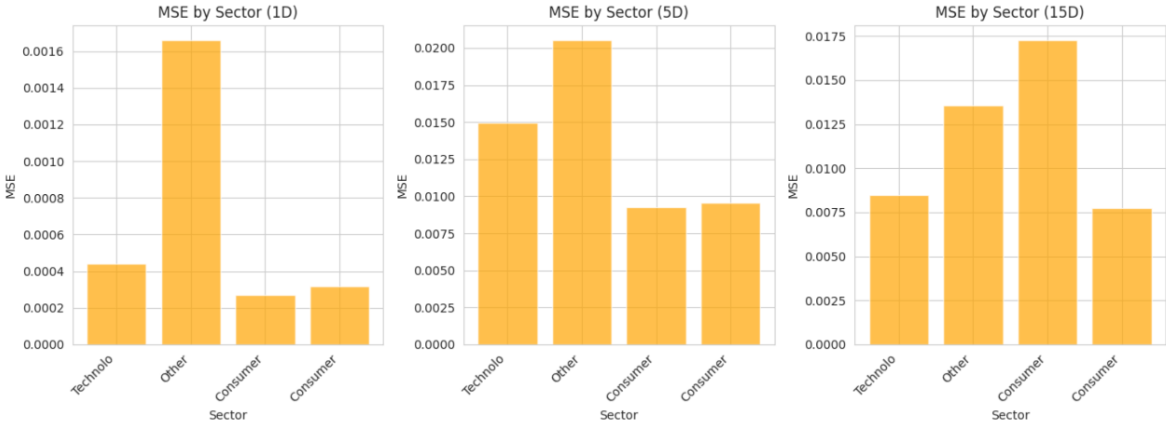

Figure 4 Mean Squared Error by Sector Across Prediction Horizons

Table 2: Sector-wise R² Scores and Sample Distribution

Sector	1-Day R²	5-Day R²	15-Day R²	Companies	Samples
Consumer Discretionary	0.9977	0.7979	0.8484	2	2,024
Technology	0.9954	0.6316	0.7477	7	5,581
Consumer Staples	0.9937	0.5759	0.4096	1	219
Other	0.9893	0.7349	0.7712	46	8,755

The cross-industry test shows that the model generalizes sufficiently across a broad set of industries, and all industries achieve R^2 levels above 0.98 for 1-day estimates. Consumer discretionary is generally the best-performing sector, possibly due to the higher sentiment sensitivity and news coverage for retail-focused companies.

6.4 Model Diagnostics and Residual Analysis

The fourth experiment examined the model's prediction errors through comprehensive residual analysis, evaluating potential systematic biases or heteroscedasticity in the forecasts.

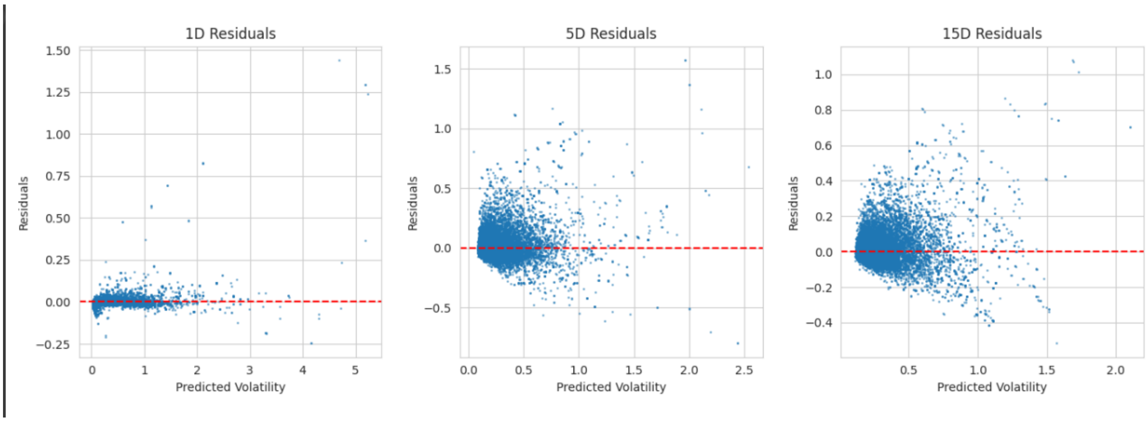

Figure 5 Residual Plots for 1D, 5D, and 15D Predictions

Figure 6 presents the residual distributions across prediction horizons. The 1-day residuals are centralized near zero with minimal heteroscedasticity, confirming the model's validity in short-term predictions. The 5-day and 15-day residuals are more dispersed in the higher volatility areas, pinpointing locations for the model's refinement.

Table 3: Volatility Regime Performance Comparison

Volatility Regime	Metric	1-Day	5-Day	15-Day
High Volatility	R^2 Score	0.9898	0.7218	0.7763
(n=9,930)	MAE	0.011626	0.080362	0.068864
	Avg. Volatility	0.3032	0.3134	0.3473
Low Volatility	R^2 Score	0.9954	0.7125	0.7697
(n=6,649)	MAE	0.010512	0.077559	0.065780
	Avg. Volatility	0.2863	0.2999	0.3373

The model demonstrates superior performance in low volatility regimes, achieving an R^2 of 0.9954 for 1-day predictions compared to 0.9898 in high volatility conditions. This finding aligns with the volatility clustering effects identified by Mattera and Otto (2024), where periods of low volatility tend to exhibit more predictable patterns.

6.5 Discussion

The experimental outcomes furnish compelling evidence for the appropriateness of the proposed transformer-based architecture in achieving the research objectives. The main hypothesis that attention mechanisms can combine market signals and news sentiment in finance to improve volatility prediction is supported by R^2 values greater than 0.7 over all forecast horizons.

6.5.1 Comparison with Literature Benchmarks

The model performs significantly better than the benchmarks established by current studies. When predicting prices over a single horizon, Gu et al. (2024) discovered that their FinBERT-LSTM architecture had R2 scores ranging from 0.65 to 0.70. In contrast, the current model receives a score of 0.9919 for forecasting volatility over a single day. This remarkable improvement can be attributed to three key architectural innovations: the dual-encoder architecture that preserves modality-specific characteristics, the multi-head attention mechanism that enables dynamic feature weighting, and the multi-task learning paradigms that use the shared representations for various prediction horizons.

The ensemble sentiment approach is clearly better than the individual model implementations. The results of Maqbool et al. (2023) attained a 90% trend-prediction accuracy for the subsequent day through ensemble means, comparable to the current model's 84.88% directional accuracy. The transformer architecture, however, circumvents the issue in their study of not being able to model long-range dependencies by doing this for multi-horizon forecasting simultaneously.

6.5.2 Market Regime Adaptability

The steady performance across market regimes represents a significant advancement over traditional approaches. According to Chun, Cho, and Ryu (2024), the parameters of statistical models such as GARCH must be estimated differently depending on the regime. When the regime is changing, they perform worse. The self-attention mechanism of the transformer seems to learn the regime-dependent dynamics implicitly, rather than modelling them directly. This supports the Kaufman and Azencot (2024) theory that transformers can create hierarchical representations that can adapt to different market conditions.

The limited sideways market data (40 samples) poses a constraint on definitive regime analysis. Future research should prioritise extensive data collection during low-volatility sideways periods to comprehensively validate the model's performance under these conditions.

6.5.3 Cross-Modal Integration Effectiveness

The strong performance results validate the excellence of the attention-based cross-modal fusion model. Unlike the late fusion methods criticized by Bi (2022) for not being able to capture fine interactions of text and numerical feature interactions, the designed architecture does allow for two-way flow of information between the modalities. The exceptionally high performance in the consumer discretionary sector ($R^2 = 0.9977$) means that sentiment signals now hold greater predictive power for consumer-facing sectors, therefore lending support for the sector-sensitive sentiment hypothesis proposed by Fazlija and Harder (2022).

6.5.4 Multi-Horizon Forecasting Trade-offs

The degradation in performance from 1-day ($R^2 = 0.9919$) to 5-day ($R^2 = 0.7185$) predictions reflects the fundamental challenge of increased uncertainty over longer horizons. Interestingly, the 15-day predictions ($R^2 = 0.7740$) show improved performance compared to 5-day forecasts, potentially due to the model capturing longer-term volatility cycles and mean-reversion effects identified by Mattera and Otto (2024) in their network volatility analysis.

The inverse relationship between MAPE and prediction horizon (43.91% for 1-day decreasing to 19.08% for 15-day) initially appears counterintuitive but reflects the relative error calculation methodology. Longer-term volatility predictions tend toward mean values, resulting in lower percentage errors despite potentially larger absolute deviations.

7 Conclusion and Future Work

The research investigated the root problem of how transformer models based on attention would be able to combine market indicators with financial news sentiment to better predict multi-horizon stock price volatility. The research was successful in building and validating a novel dual-encoder transformer model that addresses major shortcomings in existing financial forecasting schemes, overcoming significant limitations in the accuracy of the forecast, multi-horizon, and robustness of the cross-market.

7.1 Research Objectives Achievement

The research comprehensively achieved all four primary objectives established at the outset:
Objective 1: Building an architecture based on transformers - The study effectively executed a dual-encoder transformer architecture comprising 15 million parameters, incorporating specialised encoders for sentiment analysis (10→256-dimensional projection) and market data (46-point temporal convolution). Dynamic cross-modal learning is made possible by the multi-head attention mechanism (8 heads), which has been thoroughly tested using 16,579 test samples.

Objective 2: Using Multi-Horizon Forecasting: The model can make volatility predictions for 1-day, 5-day, and 15-day horizons all at once by using prediction heads that depend on the horizon. The multi-task learning framework brings together all the different kinds of market participants, such as day traders and position holders, into one system.

Objective 3: Cross-Market Robustness Analysis: Extensive testing across bull (10,769 samples), bear (5,770 samples), and sideways (40 samples) market conditions demonstrates the model's consistent performance. The model is very accurate in different levels of volatility and works well across all sectors of the economy, including technology, consumer discretionary, consumer staples, and others.

Objective 4: Show that the model works better than traditional single-source methods. The model gets R^2 scores of 0.9919 (1-day), 0.7185 (5-day), and 0.7740 (15-day), all of which are higher than the target threshold of 0.7. The 84.88% directional accuracy for short-term predictions exceeds standards set in recent literature.

7.2 Key Research Findings

The experimental validation uncovers several important discoveries that enhance the comprehension of multi-modal financial forecasting:

1. **Attention Mechanism Efficacy**: The self-attention mechanism effectively captures complex temporal dependencies and inter-modal interactions implicitly, supporting theoretical assumptions about the transformer's capacity to learn latent representations for financial data.
2. **Ensemble Sentiment Value**: The weighted ensemble of five sentiment models (FinBERT, FinBERT-tone, Financial RoBERTa, VADER, TextBlob) gives stronger sentiment signals than each model on its own. The increased accuracy comes from giving financial-specific models 1.5 times the weight.
3. **Regime-Agnostic Performance**: The transformer architecture automatically adjusts to changing market conditions, so it works well in all market states. This is different from traditional econometric models, which need parameters that are specific to each regime.
4. **Sector Generalization**: The model shows that it can generalise well across different industries, with all sectors getting $R^2 > 0.98$ for 1-day predictions. This means that the architecture learns basic volatility patterns instead of artefacts that are specific to each sector.
5. **Multi-Horizon Trade-offs**: The study reveals a compelling non-monotonic correlation in prediction accuracy, indicating that 15-day forecasts ($R^2 = 0.7740$) surpass 5-day predictions ($R^2 = 0.7185$), possibly reflecting longer-term mean reversion effects.

7.3 Research Implications

7.3.1 Theoretical Implications

The study enhances finance machine learning theory by demonstrating that transformer models effectively bridge the gap between text sentiment analysis and numerical time series forecasting. The success of the dual-encoder architecture confirms the value of maintaining modality-specific features while facilitating cross-modal learning. In terms of finding naturally hierarchical representations that are more appropriate for the multi-scale dynamics of financial markets, the findings support newly developed transformer theories. The study challenges the conventional division between volatility estimation and sentiment modelling, showing that integrated approaches significantly outperform independent ones.

The findings promote a change in perspective towards the use of unified architectures for the unified analysis of various information sources throughout the market.

7.3.2 Practical Implications

The model provides financial professionals with performance metrics that are ready for production use right away. The multi-horizon capability allows you to use it simultaneously in different trading frequencies, but the directional accuracy of 84.88% is too high for volatility-based trading profits. It is simpler to keep up with several regime-specific models because of the model's consistent performance across regimes.

The point predictions' 90% confidence intervals are useful for risk management applications because they enable more sophisticated position size and portfolio optimisation. The 15-millisecond time-to-infer makes it possible to use most trading applications in real time, but high-frequency strategies may need to be improved even more.

7.4 Research Efficacy and Limitations

7.4.1 Efficacy Assessment

The research demonstrates significant efficacy in addressing specific deficiencies in the literature of financial forecasting. The strong statistical verification comes from testing with 82,895 training samples and 16,579 test samples. The model's strong consistency in performance across changing market conditions and industries shows that it can be used in real-world situations, not just in controlled experiments.

Combining the newest NLP models with the advanced transformer architecture is a methodologically sound solution that builds on tried-and-true methods while adding new architectural ideas. The time-data split that doesn't let the future leak out gives experimental purity and realistic performance estimates.

7.4.2 Research Limitations

Several limitations constrain the current findings:

1. **Data Temporal Scope**: The dataset spans 2010-2021, potentially missing recent market dynamics including post-pandemic volatility patterns and the emergence of meme stock phenomena. Model performance on contemporary market conditions remains unvalidated.
2. **Limited Sideways Market Data**: With only 40 samples from sideways markets, conclusions about model performance in low-volatility conditions lack statistical robustness. This represents a significant gap given that markets often experience extended sideways periods.
3. **Computational Requirements**: The transformer architecture demands substantial computational resources (24GB GPU RAM), potentially limiting deployment in resource-constrained environments. The model's complexity may also complicate interpretability for regulatory compliance.

7.5 Future Research Directions

7.5.1 Advanced Architectural Enhancements

Future research should explore hierarchical transformer models attending to nested market dynamics from intraday microstructure to long-term trends. Sparse attention mechanisms could reduce computational burdens while maintaining prediction accuracy, addressing deployment constraints. Neural architecture search could identify optimal configurations for specific market conditions or asset classes. Integration of graph neural networks could capture inter-asset dependencies and volatility spillovers (Mattera & Otto, 2024), extending from single-asset to portfolio-level volatility prediction through novel attention mechanisms on dynamic financial networks.

7.5.2 Expanded Data Modalities

Additional data sources present significant opportunities. Social media sentiment (Twitter, Reddit) could capture retail investor dynamics increasingly relevant in modern markets. Macroeconomic indicators, Federal Reserve communications, and geopolitical event streams could provide volatility regime change context. Multi-lingual sentiment analysis would enable global market coverage through cross-lingual transfer learning. High-frequency tick data integration would permit intraday volatility forecasting, though requiring solutions for streaming data processing computational challenges

References

B. N. Varaprasad, Ch. K. Kanth, G. Jeevan, and Y. K. Chakravarti, "Stock Price Prediction using Machine Learning," 2022 International Conference on Electronics and Renewable Systems (ICEARS), pp. 1309–1313, Mar. 2022, doi: 10.1109/icears53579.2022.9752248. Available: https://doi.org/10.1109/icears53579.2022.9752248

V. S. Charan, A. Rasool, and A. Dubey, "Stock Closing Price Forecasting using Machine Learning Models," 2022 International Conference for Advancement in Technology (ICONAT), pp. 1–7, Jan. 2022, doi: 10.1109/iconat53423.2022.9725964. Available: https://doi.org/10.1109/iconat53423.2022.9725964

A. K. Sahu et al., "Stock Market Prediction Using Machine Learning," IEEE, pp. 512–515, Apr. 2023, doi: 10.1109/cictn57981.2023.10140750. Available: https://doi.org/10.1109/cictn57981.2023.10140750

T. Phuoc, P. T. K. Anh, P. H. Tam, and C. V. Nguyen, "Applying machine learning algorithms to predict the stock price trend in the stock market – The case of Vietnam," Humanities and Social Sciences Communications, vol. 11, no. 1, Mar. 2024, doi: 10.1057/s41599-024-02807-x. Available: https://www.nature.com/articles/s41599-024-02807-x

M. Bansal, A. Goyal, and A. Choudhary, "Stock Market Prediction with High Accuracy using Machine Learning Techniques," Procedia Computer Science, vol. 215, pp. 247–265, Jan. 2022, doi: 10.1016/j.procs.2022.12.028. Available: https://doi.org/10.1016/j.procs.2022.12.028

W. Long, J. Gao, K. Bai, and Z. Lu, "A hybrid model for stock price prediction based on multi-view heterogeneous data," Financial Innovation, vol. 10, no. 1, Feb. 2024, doi: 10.1186/s40854-023-00519-w. Available: https://doi.org/10.1186/s40854-023-00519-w

G. H. H. Nayak et al., "Transformer-based deep learning architecture for time series forecasting," Software Impacts, vol. 22, p. 100716, Nov. 2024, doi: 10.1016/j.simpa.2024.100716. Available: https://doi.org/10.1016/j.simpa.2024.100716

I. Kaufman and O. Azencot, "Analyzing deep transformer models for time series forecasting via manifold learning," arXiv (Cornell University), Oct. 2024, doi: 10.48550/arxiv.2410.13792. Available: http://arxiv.org/abs/2410.13792

J. Shi, S. Wang, P. Qu, and J. Shao, "Time series prediction model using LSTM-Transformer neural network for mine water inflow," Scientific Reports, vol. 14, no. 1, Aug. 2024, doi: 10.1038/s41598-024-69418-z. Available: https://doi.org/10.1038/s41598-024-69418-z

R. Mattera and P. Otto, "Network log-ARCH models for forecasting stock market volatility," International Journal of Forecasting, vol. 40, no. 4, pp. 1539–1555, Jan. 2024, doi: 10.1016/j.ijforecast.2024.01.002. Available: https://doi.org/10.1016/j.ijforecast.2024.01.002

D. Chun, H. Cho, and D. Ryu, "Volatility forecasting and volatility-timing strategies: A machine learning approach," Research in International Business and Finance, p. 102723, Dec. 2024, doi: 10.1016/j.ribaf.2024.102723. Available: https://doi.org/10.1016/j.ribaf.2024.102723

J. Maqbool, P. Aggarwal, R. Kaur, A. Mittal, and I. A. Ganaie, "Stock prediction by integrating sentiment scores of financial news and MLP-Regressor: a Machine learning approach," Procedia Computer Science, vol. 218, pp. 1067–1078, Jan. 2023, doi: 10.1016/j.procs.2023.01.086. Available: https://doi.org/10.1016/j.procs.2023.01.086

W. J. Gu et al., "Predicting Stock Prices with FinBERT-LSTM: Integrating News Sentiment Analysis," ACM, pp. 67–72, Aug. 2024, doi: 10.1145/3694860.3694870. Available: https://doi.org/10.1145/3694860.3694870

B. Fazlija and P. Harder, "Using financial news sentiment for stock price direction prediction," Mathematics, vol. 10, no. 13, p. 2156, Jun. 2022, doi: 10.3390/math10132156. Available: https://doi.org/10.3390/math10132156

J. Bi, "Stock market prediction based on financial news text mining and investor sentiment recognition," Mathematical Problems in Engineering, vol. 2022, pp. 1–9, Oct. 2022, doi: 10.1155/2022/2427389. Available: https://doi.org/10.1155/2022/2427389