TCNAttention-RAG: Stock Prediction and Fraud Detection Framework Based on Financial Report Analysis

Yuan Liu¹, Yilong Liu¹, Ziang Qi², Yao Xiao³, Xiaoyan Guo^{1*}

¹School of Mathematics, Northwest University, Xi'an, China
²Duke University, Durham, USA
³University of Southern California, Los Angeles, USA
*Corresponding Author. Email: guoxiaoyan617@nwu.edu.cn

Abstract: Due to the high volatility of financial markets and the prevalence of financial fraud, real-time stock market forecasting for listed companies remains a challenging task. To address these challenges, this study proposes TCNAttention-RAG, a hybrid deep learning framework integrating Temporal Convolutional Network (TCN), Multi-Layer Perceptron (MLP), Attention Mechanism, and Retrieval-Augmented Generation (RAG) for enhanced stock price forecasting. The model leverages TCN for temporal feature extraction, MLP for nonlinear representation, and Attention for feature weighting, while RAG dynamically retrieves key financial insights from corporate reports to improve predictive accuracy. Using NASDAQ-listed stock price data (2014–2020), combined with corporate financial reports, market transaction data, and macroeconomic indicators, a multi-dimensional dataset is constructed. Experimental results demonstrate that TCNAttention-RAG outperforms traditional models in accuracy and recall, effectively capturing stock price fluctuations. Despite its limitations in handling extreme market events, the model exhibits high reliability and predictive robustness. This study introduces a multi-modal data-driven approach to financial forecasting, offering insights into intelligent financial analysis and enhancing decision-making in volatile markets.

Keywords: Stock Prediction, TCN, MLP, Attention Mechanism, RAG

1. Introduction

The stock market plays a pivotal role in the global financial system, with the New York stock exchange exchange serving as a hub for technology and financial enterprises. Stock prices fluctuate in response to macroeconomic policies, corporate financial performance, and investor sentiment, creating substantial forecasting challenges due to the market's inherent volatility and nonlinear dynamics.

Accurate stock price prediction is critical for investors, corporate managers, and financial institutions. Investors rely on predictive models to optimize trading strategies and mitigate risks, while financial institutions leverage forecasts for portfolio management. However, conventional time-series models struggle with the complexity of stock price movements, failing to effectively capture nonlinear dependencies and external financial influences. Financial fraud, such as misrepresentation of company performance or concealment of financial risks in corporate reports, can severely distort market information, mislead investors, and impact stock prices. Traditional financial

analysis methods often fail to effectively identify these fraudulent activities, increasing investment risks.

To address these limitations, this study proposes TCNAttention-RAG, a hybrid deep learning framework that integrates Temporal Convolutional Networks (TCN) for time-series feature extraction, Multilayer Perceptron (MLP) for nonlinear representation, and Attention Mechanism for dynamic feature weighting. Additionally, Retrieval-Augmented Generation (RAG) is incorporated to enhance financial report analysis by retrieving and synthesizing key textual information, thereby enriching the model's predictive insights.

Using historical stock price data (2014–2020) from New York stock exchange-listed companies, along with corporate financial reports, market transaction data, and macroeconomic indicators, this study constructs a multi-dimensional dataset to improve predictive accuracy. Experimental results confirm that TCNAttention-RAG surpasses traditional forecasting models across key performance metrics. By integrating structured financial indicators with unstructured text from financial disclosures, this research contributes a novel approach to stock prediction and intelligent financial analytics.

2. Literature review

Stock price prediction has been extensively studied using machine learning and deep learning techniques. Prior research can be categorized into three main areas: traditional machine learning approaches, deep learning models, and retrieval-augmented methods for financial text analysis.

Early studies leveraged classical machine learning models such as Support Vector Machines (SVM), Random Forest (RF), and LASSO regression to predict stock prices. Mintarya et al. [1] reviewed 30 studies on machine learning in stock market prediction, concluding that while neural networks were predominant, SVM and RF also demonstrated effectiveness in specific market conditions. Gao et al. [3] further explored feature selection techniques like LASSO and Principal Component Analysis (PCA), demonstrating that LASSO-based models outperformed PCA in reducing dimensionality while preserving predictive accuracy. However, these methods often struggled with nonlinear market dynamics and failed to incorporate unstructured financial data.

The introduction of Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks significantly enhanced stock prediction by capturing temporal dependencies. Lu et al. [2] proposed a CNN-BiLSTM-Attention (AM) model, where CNN extracted data features, BiLSTM processed sequential dependencies, and attention mechanisms emphasized critical historical patterns. Their model achieved superior performance in terms of Mean Absolute Error (MAE) and Root Mean Square Error (RMSE), demonstrating its effectiveness in complex time-series forecasting.

Recent research has explored advanced architectures such as Temporal Convolutional Networks. Wan et al. [4] developed a Multivariate TCN for financial time-series forecasting, proving its ability to handle multivariate interactions effectively. Similarly, Dudukcu et al. [5] combined TCN with RNNs to enhance chaotic time-series predictions, providing a methodological reference for capturing nonlinear and uncertain market fluctuations.

Feature engineering and hybrid models have also gained traction. Yun et al. [6] proposed a GA-XGBoost framework with an enhanced feature selection process, achieving improved accuracy by identifying optimal input variables. Meanwhile, Hu et al. [7] surveyed deep learning applications in financial forecasting, concluding that LSTM combined with reinforcement learning exhibited strong potential in adaptive financial modeling.

The integration of text-based financial data into predictive models has been an emerging research direction. Zhang et al. [8] introduced a CoATSMP model, which adaptively fused numerical stock data with textual financial reports using attention-based fusion techniques. Their results highlighted that models incorporating both structured price data and textual financial insights outperformed

traditional price-only models, with textual sentiment playing a crucial role in market movement predictions.

Further, Retrieval-Augmented Generation has been explored for financial text processing. Cuconasu et al. [9] demonstrated how RAG can enhance knowledge retrieval by extracting key financial insights from large-scale documents. Their research underscored the advantage of integrating domain-specific retrieval into deep learning models, significantly improving the interpretability of financial predictions.

3. Data introduction

To construct a robust predictive framework, this study utilizes structured financial data and unstructured financial reports, integrating multiple data sources to enhance stock price forecasting. This study utilizes the NYSE Financial Fundamentals dataset (2014–2020), including financial indicators like revenue, EPS, and net income, along with the AXA stock trading dataset (2014–2020) containing market variables such as prices and trading volume. The Three-Sigma Rule and mean imputation ensure data integrity. Techniques like 30-day moving averages, log returns, and correlation analysis identify key features. PCA is applied for dimensionality reduction, optimizing the feature set and enhancing model stability for stock price prediction.

Figure 1 illustrates the stock price trends of AXA from 2014 to 2020, reflecting the dynamic changes in the company's stock price over time. The chart highlights the volatility patterns of the stock price, which may be attributed to factors such as seasonal market fluctuations, corporate performance, macroeconomic conditions, and shifts in investor sentiment. By analyzing this chart, we can gain a comprehensive understanding of the overall trajectory of AXA's stock price during this period, offering valuable insights for market trend analysis and informed investment decision-making.



Figure 1: Overall stock price changes of AXA company

4. TCNAttention-RAG model

The core module of this study is based on the TCNAttention-RAG framework, with the goal of predicting the stock prices of New York stock exchange AXA listed companies from 2014 to 2020. It combines structured financial indicators with multimodal data from unstructured financial report texts for analysis.

The TCNAttention-RAG model architecture mainly consists of the following three parts. Figure 2 shows the structural framework of this model.

Proceedings of the 7th International Conference on Computing and Data Science DOI: 10.54254/2755-2721/2025.22702



Figure 2: Model structure

4.1. Time convolutional network (TCN)

The TCN component serves as the primary temporal feature extractor, using dilated causal convolution to process sequence data. The TCN structure contains multiple convolutional blocks with exponentially increasing dilation rates [10]. For a given input sequence x, the TCN operation can be expressed as $TCN(x) = F(x * f_d)$, where f_d represents the dilated convolution filter and F is the activation function.



Figure 3: Time convolutional network

Figure 3 depicts the structure of a Temporal Convolutional Network (TCN). Panel (a) illustrates a TCN layer where input sequences are processed through dilated causal convolutions at varying dilation rates (d=1, d=2, etc.), allowing the model to capture long-range dependencies. Panel (b) shows the TCN residual block, starting with dilated causal convolution for past data causality, followed by a weight normalization layer for faster convergence, and a ReLU activation layer for non-linearity. A dropout layer ensures better generalization. Residual connections improve gradient flow and prevent vanishing gradients, while a 1x1 convolution aligns input and output dimensions for residual addition. This architecture aids in learning complex patterns without deeper stacks, essential for tasks like stock price prediction.

4.2. Attention mechanism

The attention mechanism enhances the model's ability to focus on relevant temporal patterns by assigning different weights to different time steps. The attention weights are calculated as $\alpha_t = \text{softmax}(v^T \tanh(Wh_t + b))$, where h_t represents the hidden state at time t, W is a learnable weight matrix, v is the context vector, and b is the bias term. The final context vector c is calculated as $c = \sum \alpha_t h_t$.

The complete forward propagation of the model can be expressed as:

$y = f_{out}(Attention(MLP(TCN(x))))$

where f_{out} represents the final output layer that generates price predictions.

4.3. RAG enhanced processing and querying of financial report text

In addition, the Retrieval Augmented Generation technology is introduced to extract key paragraphs (such as "Management Discussion and Analysis" and "Risk Factors") from AXA's annual financial reports (such as 10-K and 10-Q reports) and construct a domain knowledge base. By pre training language models (such as BERT) to semantically encode text, generate context relevant vector representations, and combine similarity retrieval to match historical text fragments, qualitative information that is not explicitly included in structured data (such as financial report information, strategic adjustments, policy impacts) is supplemented, thereby enhancing the explanatory power and predictive robustness of the model.

5. Model results analysis

5.1. Analysis of model comparison results

The integration of Retrieval-Augmented Generation significantly enhances the model's ability to acquire domain-specific financial knowledge, thereby improving its capacity for financial data analysis and stock price prediction. Specifically, RAG enables a financial report information retrieval module, which dynamically integrates structured financial indicators (e.g., revenue, earnings per share (EPS), debt-to-equity ratio) with unstructured textual data (e.g., Management Discussion and Analysis (MD&A), company announcements). This hybrid approach allows the model to retrieve and incorporate relevant financial insights, reducing prediction errors caused by knowledge limitations.

Experimental results confirm that RAG consistently improves prediction accuracy and recall rates across various model configurations, reinforcing its essential role in financial forecasting. By enhancing the model's comprehension of financial reports, RAG reduces predictive uncertainty and provides a robust foundation for stock price forecasting models.

Table 1 summarizes the experimental results of different model configurations, comparing the accuracy and recall rates among baseline models (gpt-3.5-turbo and gpt-4.0) and their RAG-enhanced variants, highlighting the significant performance improvements achieved through RAG technology integration.

Model	Accuracy rate	Recall	
gpt-3.5-turbo	50.7%	73.3%	
gpt-3.5-turbo+RAG	65.8%	89.2%	
gpt-3.5-turbo-1106	60.4%	82.3%	
gpt-3.5-turbo-1106+RAG	79.5%	90.1%	
gpt-4.0	80.2%	92.7%	

Fable	1.	Model	results
auto	1.	IVIUUUI	results

5.2. Model prediction results

This study conducts an in-depth analysis of the performance of the TCNAttention-RAG model in the task of predicting the stock prices of AXA in the New York stock exchange market from 2014 to 2020. Through a comprehensive examination of visual representations and quantitative evaluation

metrics, the robustness and predictive capabilities of the model under different market conditions are validated.



Figure 4: Training loss and validation loss

Figure 4 results illustrates the model's effectiveness in capturing both macro-level trends and micro-level price movements.



Figure 5: AXA stock price prediction results

Figure 5 shows the prediction results of the stock price of AXA (2014-2020), a New York stock exchange listed company, based on the TCN-MLP Attention RAG framework. The blue curve in the figure represents the actual stock price trend, while the orange and green curves represent the predicted values of the model on the training and testing sets, respectively. The horizontal axis represents the time series (0-1400 units), and the vertical axis represents the stock price range (16-28 US dollars). The experimental results show that the predicted curve of the model is highly consistent with the actual price, especially during the fluctuation cycle from 2017 to 2019. The test set prediction (green) can accurately capture the periodic peaks and pullback trends of the stock price, indicating that the model has strong temporal pattern recognition ability.

6. Conclusions

This study proposes the TCN-MLP-Attention-RAG framework, which integrates Temporal Convolutional Networks for sequential modeling, Multilayer Perceptron for nonlinear feature transformation, and attention mechanisms for dynamic weighting, effectively capturing complex stock price fluctuations. Additionally, Retrieval-Augmented Generation enhances the model by extracting semantic insights from financial reports (e.g., 10-K and 10-Q filings), compensating for the limitations of structured financial indicators. Experimental results demonstrate that RAG-driven financial report analysis significantly reduces prediction errors, particularly during earnings release periods, validating the effectiveness of multimodal data fusion.

Despite its superior predictive performance, the framework has certain limitations: 1) Delayed response to extreme market events (e.g., black swan events) due to external knowledge base update

latency; 2) Suboptimal prediction accuracy for peak-valley fluctuations in highly volatile short-term movements.

By bridging structured financial data with unstructured textual insights, this study establishes an extensible framework for financial time-series prediction, providing investors with data-driven decision support tools. These findings underscore the potential of intelligent financial analysis systems in complex market environments, offering both theoretical foundations and practical applications for future research.

Acknowledgements

This work was supported by the Shaanxi Basic Science Research Institute Scientific Research Project (No. 22JSZ008).

References

- [1] Mintarya, L. N., Halim, J. N., Angie, C., Achmad, S., & Kurniawan, A. (2023). Machine learning approaches in stock market prediction: A systematic literature review. Procedia Computer Science, 216, 96-102.
- [2] Lu, W., Li, J., Wang, J., & Qin, L. (2021). A CNN-BiLSTM-AM method for stock price prediction. Neural Computing and Applications, 33(10), 4741-4753.
- [3] Gao, Y., Wang, R., & Zhou, E. (2021). Stock prediction based on optimized LSTM and GRU models. Scientific Programming, 2021(1), 4055281.
- [4] Wan R, Mei S, Wang J, et al. Multivariate temporal convolutional network: A deep neural networks approach for multivariate time series forecasting[J]. Electronics, 2019, 8(8): 876.
- [5] Dudukcu H V, Taskiran M, Taskiran Z G C, et al. Temporal Convolutional Networks with RNN approach for chaotic time series prediction[J]. Applied soft computing, 2023, 133: 109945.
- [6] Yun, K. K., Yoon, S. W., & Won, D. (2021). Prediction of stock price direction using a hybrid GA-XGBoost algorithm with a three-stage feature engineering process. Expert Systems with Applications, 186, 115716.
- [7] Hu, Z., Zhao, Y., & Khushi, M. (2021). A survey of forex and stock price prediction using deep learning. Applied System Innovation, 4(1), 9.
- [8] Zhang, Q., Zhang, Y., Bao, F., Liu, Y., Zhang, C., & Liu, P. (2024). Incorporating stock prices and text for stock movement prediction based on information fusion. Engineering Applications of Artificial Intelligence, 127, 107377.
- [9] Cuconasu, F., Trappolini, G., Siciliano, F., Filice, S., Campagnano, C., Maarek, Y., ... & Silvestri, F. (2024, July). The power of noise: Redefining retrieval for rag systems. In Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (pp. 719-729).