ISSN 2409-2665
Journal of Logistics, Informatics and Service Science
Vol. 12 (2025) No. 2, pp. 256-269
DOI:10.33168/JLISS.2025.0216

Differential Event Impact Analysis for Stock Prediction Using Heterogeneous Graph Attention Networks

Zuxin Wang [1], Md Yasin Ida [1], Zuhao Jin [2] and Abdul Wahab Noor Maimun [1]

[1] Putra Business School, Universiti Putra Malaysia, Selangor,43400, Malaysia

[2] College of Computer Science and Technology, Zhejiang University of Technology, Hangzhou, 310000, China

pbs23204127@grad.putrabs.edu.my(Corresponding author); ida@putrabs.edu.my; jzh19990628@gmail.com; noormaimun@putrabs.edu.my

Abstract. Stock market prediction remains a challenging task due to the complex and dynamic interactions among various influencing factors. This paper proposes EDHAN (Event-Driven Heterogeneous Attention Network), a novel model designed to address two key limitations in existing methods: the varying impact of the same event on different stocks and the cross-stock transmission of event effects. EDHAN constructs a heterogeneous graph that integrates technical indicators, fundamental data, and financial news, enabling a holistic representation of market dynamics. A dual-level attention mechanism—comprising node-level and semantic-level attention—allows the model to capture fine-grained dependencies between events and stock movements. Experiments on stocks from the Shanghai and Shenzhen 300 Index (2020–2024) show that EDHAN outperforms state-of-the-art baselines, improving prediction accuracy by 28% and investment returns by 32.15%. Furthermore, EDHAN demonstrates strong risk-adjusted performance, achieving a Sharpe ratio of 3.4 and a maximum drawdown of only 6.4%, underscoring its practical value for real-world investment decision-making.

Keywords: Stock trend prediction; Heterogeneous graph; Event-driven analysis

1. Introduction

Predicting stock market trends can assist investors in understanding market movements and patterns, thereby reducing investment risks and securing considerable profits (Fama, 1995). However, stock market volatility has intensified due to an increasing frequency of macroeconomic shocks, geopolitical events, and policy interventions (Bekaert et al., 2013; Baker et al., 2016; Reboredo et al., 2017). These factors can significantly impact stock prices, contributing to market volatility and uncertainty. Among these factors, financial news events often serve as immediate triggers for market fluctuations (Tetlock, 2007). Events such as mergers and acquisitions, regulatory changes, or earnings reports can cause sharp stock movements, both within specific companies and across related sectors. Therefore, incorporating such event-level information into stock prediction models is essential for capturing real-world market dynamics and improving forecasting performance. In particular, event-driven stock prediction has gained attention. Fig. 1 illustrates the process of how financial news events impact the stock market.

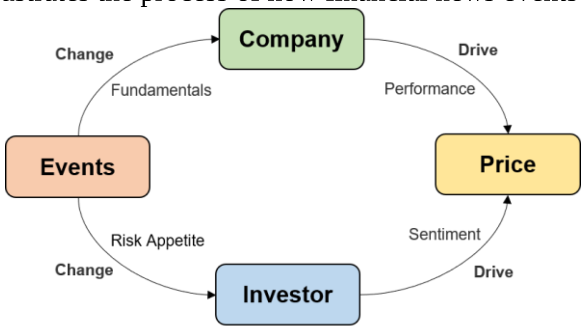

Fig 1: The Process of How Financial News Events Impact the Stock Market.

Although various event-driven stock prediction models have been proposed, leveraging text mining, knowledge graphs, and deep learning (Sun et al. 2022; Wu et al. 2023), they mainly overlook two issues. Firstly, they ignore the differences in the impact of the same event on different stocks. For example, Geely's acquisition of Volvo led to a significant increase in its stock price, while Tsinghua Unigroup's aggressive acquisitions resulted in the company's bankruptcy restructuring. This indicates that factors such as company strategy, management capabilities, and market environment are equally important for financial event analysis. Secondly, these models overlook the transmission of event impacts among related stocks. For instance, Alibaba's initial public offering boosted the development of internet stocks, demonstrating that the surge in a company's business can indicate the development trend of its industry, thereby affecting the performance of related stocks.

To address these gaps, we propose a multi-factor event-driven stock prediction model based on a heterogeneous graph attention network (EDHAN). This model analyzes financial news events to assess their impact on stock trends while incorporating technical and fundamental data as contextual information. This integration allows for a deeper investigation into how the same event affects stocks under different market conditions. Furthermore, we represent these multi-source data as heterogeneous graph-structured data and process them using a heterogeneous graph attention network to explore the transmission effects of financial events among related stocks. Finally, we utilize the embedded representations of graph nodes to predict future stock trends. In summary, the EDHAN model is capable of real-time processing and analysis of large-scale data, thereby providing more accurate stock price predictions and assisting investors in making wiser decisions. The core advantage of the model lies in its integration of multiple data sources and event-driven approach, enabling it to demonstrate stronger adaptability and predictive power in the face of market volatility and uncertainty. The main contributions of this work can be concluded as follows:

- We propose a novel data fusion framework that transforms multi-source stock market information including technical indicators, fundamental data, and financial news into a unified heterogeneous graph structure, enabling richer market context modeling.
- Based on this framework, we design EDHAN, an event-driven heterogeneous graph attention

network capable of performing real-time and multi-event prediction using fused information sources.

- We conducted a series of comparative experiments and investment simulation experiments to validate the effectiveness of our model. And the results indicate our model consistently outperforms baseline models across multiple metrics, demonstrating its superior performance.

2. Related Work

2.1. Event-Driven Stock Prediction

Event-driven stock prediction models differ from traditional time-series models (Yang et al. 2018; Naik et al. 2019; Lu et al. 2024) by incorporating external events or related information, rather than relying solely on historical prices. These events may include significant occurrences such as company financial reports, product launches, mergers and acquisitions, and changes in industry policies that have a substantial impact on stock prices. In recent years, with the widespread adoption and development of the internet and mobile terminals, such information is often disseminated to investors in the form of financial news. Investors make individual decisions based on the market signals conveyed by financial news events, thereby influencing stock trends. Hence, many researchers attempt to mine financial news event information to predict stock trends.

(Ding et al. 2014) employed Open Information Extraction (OpenIE) technology to extract structured news events from large-scale public news sources, aiming to predict stock price movements based on these events. (Zou et al. 2022) proposed a method called SRLP, which uses Semantic Role Labeling (SRL) to create embedding representations for each news paragraph. Transformer is then employed to process these embedding vectors for predicting stock movement trends. (Jiang et al. 2022) propose a multi-task learning framework based on BERT that simultaneously predicts stock trends and classifies financial events, enabling better alignment between event semantics and market reactions. (Yang et al. 2023) incorporate causal inference techniques to model the impact of financial news events on stock returns, aiming to filter out spurious correlations and identify truly influential events. (Ma et al. 2023) proposes a Multi-source Aggregated Classification (MAC) model that integrates numerical features, market-driven news sentiments of target stocks, and news sentiments of related stocks, incorporating pre-trained sentiment embeddings and a graph convolutional network to enhance stock price movement prediction accuracy and trading performance. (Liu et al. 2024) proposes a novel Multiscale Multimodal Dynamic Graph Convolution Network (Melody-GCN) that integrates numerical and textual features, refines temporal information through a multiscale architecture, and captures dynamic spatio-temporal stock relations, achieving superior performance in stock price prediction and trading simulations.

Overall, current event-driven stock prediction methods often start with news event text to explore the impact of financial events on the stock market, aiming to enhance the model's adaptability to sudden events in the stock market. However, existing methods often overlook two practical issues: the differences in events' effects on different stocks and the interdependence among stocks. Based on this, this paper proposes a more comprehensive financial event analysis method building upon previous research.

2.2. Heterogeneous Graphs

Heterogeneous graphs are structured graphs that consist of multiple types of nodes and edges, corresponding to complex relationships between different entity types. Due to these structural complexities, traditional neural network models such as Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and Graph Convolutional Networks (GCNs) struggle to effectively process heterogeneous graph data (Yang et al. 2023; Yang et al. 2024). This limitation has led to a growing research focus on heterogeneous graph neural network models, which is designed to capture and learn the diverse relationships within these graphs.

(Tan et al. 2022) proposes FinHGNN, a conditional heterogeneous graph neural network that efficiently captures multiple spillover effects in asset pricing by preserving relational attributes and using a conditional message-passing mechanism. (Wu et al. 2022) introduce a meta-path enhanced heterogeneous graph attention network that leverages semantic information from multiple meta-paths to improve recommendation performance. (Wang et al. 2023) introduces HGISD, a novel negative sample-free self-supervised learning framework for heterogeneous graph tasks, utilizing iterative similarity distillation (IMCD and IMSD) to preserve inter-meta path correlation and intra-meta path locality. (Ceskoutsé et al. 2024) proposes HeteroKGRep, a novel drug repositioning model that leverages heterogeneous biomedical knowledge graphs to enhance representation learning, achieving state-of-the-art performance in predicting repurposing opportunities and establishing a new paradigm for knowledge-guided drug discovery. (Zhou et al. 2024) propose a heterogeneous graph contrastive learning framework with generative augmentation, which enhances representation learning by generating diverse views to better capture structural and semantic information. (Shi et al. 2025) proposes JLR-GCN, a joint label-aware and relation-aware graph convolutional network that enhances heterogeneous graph representation learning by leveraging label attributes and structural information.

Currently, heterogeneous graphs have been widely used in fields such as recommendation systems, social network analysis, and knowledge graphs, proving useful in describing relationships in complex systems and networks. However, their application in the financial industry is still at a nascent stage.

3. Methods

3.1. Overall Framework

The main process of the EDHAN model consists of three parts: graph representation, graph learning, and stock trend prediction. In the graph representation step, we integrate news, fundamental, and technical information from the stock market into heterogeneous structured graph data based on financial events. During the graph learning process, we employ a heterogeneous graph attention network to obtain embeddings for stock nodes. Finally, we predict the movement trends of corresponding stocks based on these node embeddings. Fig. 2 illustrates the framework of our model, which will be elaborated on in sections 3.2-3.4.

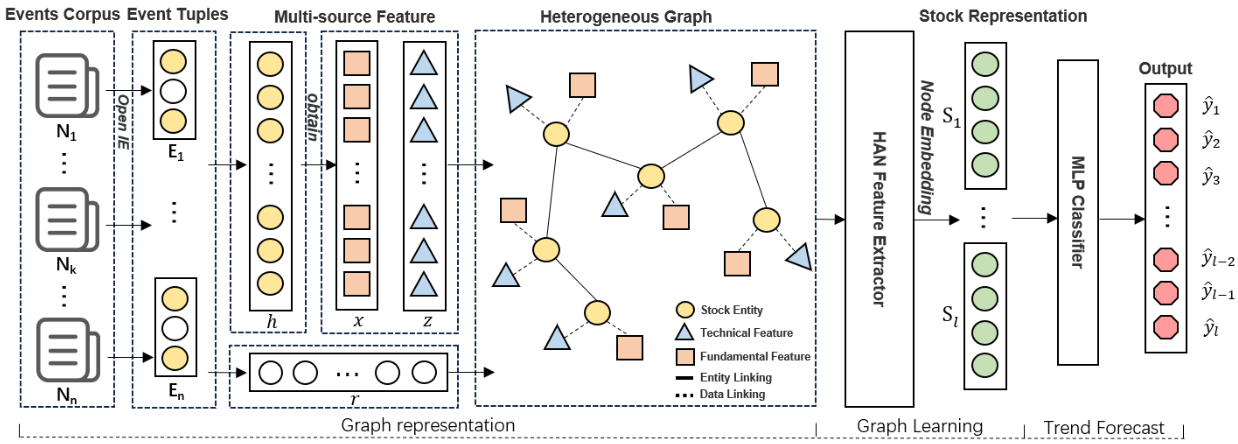

Fig 2: Main Framework of the EDHAN Model

3.2. Graph Representation

In the graph representation process, the news event corpus $D = (N_1, N_2, \cdots, N_n)$ is used as input, where N_i represents the i-th news event text segment. Open Information Extraction (Etzioni et al. 2008) is employed in this work to extract relational triples from financial news articles, providing a lightweight and flexible means to convert unstructured text into structured representations. While OpenIE may introduce noise due to general-purpose parsing, its impact is mitigated in our framework for two reasons. First, we apply domain-specific post-processing, filtering out tuples with generic or non-financial predicates and retaining only relations involving identified stock entities or sector-specific keywords. Second, as our downstream graph-based model aggregates information across multiple events and entities, individual tuple-level noise is diluted, and important structural patterns can still be effectively captured. Therefore, despite its limitations, OpenIE remains a practical and efficient choice for large-scale event extraction in the financial news domain. Then, each event is structured into an event tuple $E_i = (h_{2i-1}, r, h_{2i})$, where r represents the event type such as acquisition, merger, sale, etc., and h_i represents the stock entity. Subsequently, based on different stock entities h_i, we obtain the corresponding technical feature x_i and fundamental feature z_i for each stock. The specific features used here can be found in the appendix. Then, we link the stock entity set, technical feature set, and fundamental feature set into a heterogeneous graph $HG = (V, E)$, where $V = \{h_1, \dots, h_l, x_1, \dots, x_l, z_1, \dots, z_l\}$ represents the node features, E represents the edge features corresponding to data links and stock entity links, and l denotes the number of stock entities. The x_i, z_i here will be calculated with h_i as a whole node. At this point, the mapping function of node types in the heterogeneous graph is denoted as $\phi : V \rightarrow A$, and the mapping function of connection types is denoted as $\psi : E \rightarrow R$, where A and R represent the node types and connection types, respectively, and $|A| + |R| \geq 2$. Given a node i and a meta-path Φ in the heterogeneous graph, the meta-path-based neighbors N_i^{Φ} of node i are defined as the set of nodes connected to node i via the meta-path Φ. A representative meta-path in our heterogeneous graph is Stock → Event → Stock, which captures the co-occurrence of multiple stocks within the same financial event. This relation reflects scenarios where companies are jointly affected by market-moving events. For example, in the event "Apple increases chip orders from TSMC," both Apple and TSMC are linked through the meta-path, indicating potential correlated movements.

3.3. Graph Learning

In the graph learning process, we employ the Heterogeneous Graph Attention Network (HAN) structure proposed by (Wang et al. 2021) to process the heterogeneous graph HG. This network applies attention mechanisms at both the node-level and semantic-level to capture fine-grained representations. The motivation behind using these two levels of attention is to effectively handle both local and global dependencies in the graph structure, which is essential for capturing complex event-driven relationships. The node-level attention mechanism focuses on capturing local dependencies between nodes, which is crucial in understanding how individual stocks interact with their immediate neighbors. By applying attention at this level, the model can selectively prioritize certain relationships over others, ensuring that more relevant neighboring nodes contribute to the stock's representation. This allows the model to better capture short-range correlations and refine the local context of each node, which is important for

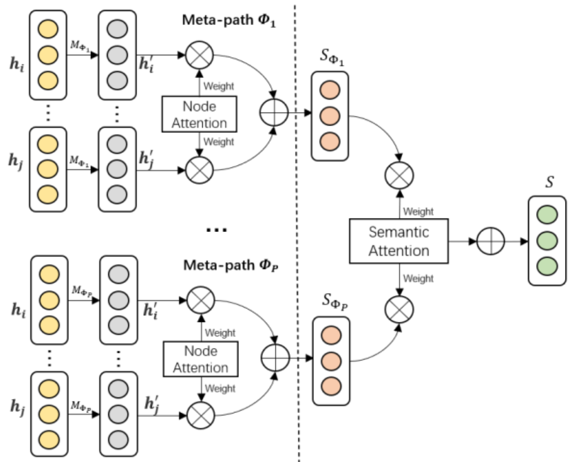

Fig 3: The Process of Graph Learning in Node-level and Semantic-level Attention

understanding immediate market reactions to specific events. The semantic-level attention mechanism complements the node-level attention by focusing on global dependencies within the graph. While node-level attention handles local interactions, semantic-level attention allows the model to weigh the significance of different types of relationships. By learning which meta-paths are more important, the model can capture the broader, higher-level context of the graph. This is crucial for understanding the global patterns that emerge from complex event-driven interactions, enabling the model to make more accurate predictions based on a holistic view of the market. The detailed calculation process is illustrated in Fig. 3.

3.3.1 Node-level Attention

Node-level attention aggregates features between nodes on different meta-paths. In this process, for each node feature h_i, there exists a transformation matrix M_{ϕ_i} that maps feature vectors of different dimensions within the node to the same dimension. The calculation is as follows:

$$h_i' = M_{\phi_i} \cdot h_i \tag{1}$$

Where h_i' is the transformed feature vector. Then, self-attention is employed to dynamically learn the weight of each node. For a given pair of nodes (i, j) connected by a meta-path Φ, the node attention value e_{ij}^{Φ} represents the importance of node j to node i, calculated as:

$$e_{ij}^{\Phi} = att_{node}\left(h_i', h_j'; \Phi\right) = \sigma(a_{\Phi}^{T} \cdot [h_i'] || \left[h_j'\right]) \tag{2}$$

Where att_{node} represents node-level attention, σ is the activation function, $||$ denotes concatenation operation, $j \in N_i^{\Phi}$, a_{Φ}^{T} represents the node-level attention vector. Then, by applying softmax operation, we obtain the weight coefficients for each feature vector:

$$\alpha_{ij}^{\Phi} = softmax(e_{ij}^{\Phi}) \tag{3}$$

Finally, the calculation of the final embedding s_i^{Φ} for node i in meta-path Φ is:

$$s_i^{\Phi} = \sigma\Big(\sum_{j \in N_i^{\Phi}} \alpha_{ij}^{\Phi} \cdot h_j'\Big) \tag{4}$$

After this node-level attention for a given set of meta-paths $\{\Phi_1, \Phi_2, \dots, \Phi_P\}$, we obtain P groups of semantically related node embeddings, denoted as $\{S_{\Phi_1}, S_{\Phi_2}, \dots, S_{\Phi_P}\}$.

3.3.2 Semantic-level Attention

Semantic-level attention takes the P group of node embeddings as input, which are obtained from node-level attention. The learning weights for each meta-path $(\beta_{\Phi_1}, \beta_{\Phi_2}, \dots, \beta_{\Phi_P})$ can be expressed as follows:

$$\left(\beta_{\Phi_1}, \beta_{\Phi_2}, \dots, \beta_{\Phi_P}\right) = att_{sem}(S_{\Phi_1}, S_{\Phi_2}, \dots, S_{\Phi_P}) \tag{5}$$

Here, $att_{\ sem}$ denotes the semantic-level attention operation, which aims to learn the significance of each meta-path. It calculates the semantic-level attention vector q and the embeddings of each edge type in order. And it is computed as:

$$\omega_{\Phi_P} = \frac{1}{|V|} \sum_{i \in V} q^T \cdot \tanh\left(W \cdot s_i^{\Phi_P} + b\right) \tag{6}$$

Where q is the semantic-level attention vector, W is the learnable weight matrix, and b is the learnable bias vector. Finally, by applying softmax to ω_{Φ_P} obtained from equation (6), the attention weights for each meta-path category can be obtained.

$$\beta_{\Phi_P} = \frac{\exp\left(\omega_{\Phi_P}\right)}{\sum_{k=1}^{P} \exp\left(\omega_{\Phi_k}\right)} \tag{7}$$

Finally, the embedding representation of the root node i across all meta-paths is:

$$S = \sum_{n=1}^{P} \beta_{\Phi_P} \cdot S_{\Phi_P} \tag{8}$$

By considering all stock nodes $\{h_1, h_2, \ldots, h_l\}$ as root nodes, we obtain all stock embeddings $\{S_1, S_2, \ldots, S_l\}$, and subsequently predict the trend of corresponding stocks based on all stock embeddings.

3.4. Stock Trend Prediction

In the stock trend prediction process, using all stock embeddings $\{S_1, S_2, \ldots, S_l\}$ as input, a Multi-Layer Perceptron (MLP) network is employed for binary classification of stock price movement (rise or fall). The specific computation process is as follows:

$$\{\hat{y}_1, \hat{y}_2, \ldots, \hat{y}_l\} = MLP(\{S_1, S_2, \ldots, S_l\}) \tag{9}$$

Here, $\hat{y}_1$ represents the output of the model, indicating the probability of stock i rising in price after the occurrence of a news event, with $\hat{y}_i \in [0, 1]$. In this task, cross-entropy loss is employed as the loss function to calculate prediction errors and optimize network parameters, defined as:

$$L = -\sum_{i=1}^{l} y_i \log(\hat{y}_i) + (1 - y_i)\log(1 - \hat{y}_1) \tag{10}$$

Where y_i is the true price movement of the ith stock (0 for decrease, 1 for increase).

4. Experimental Evaluation

4.1. Experimental Data

We selected the constituents of the Shanghai and Shenzhen 300 Index (HS300) as the research subjects and obtained financial news records from the Financial News Database (CFND) on the China National Research Data Sharing Platform. We used the financial news records from October 2020 to January 2024 as the overall dataset, where records from October 2020 to October 2022 were used as the training set, records from November 2022 to July 2023 were used as the test set, and records from August 2023 to January 2024 were used as the validation set. Basic and technical data for the corresponding stock entities were obtained online from the Company Basic Information Database (CBID) and Tushare data platform. Table 1 shows the detailed description of the news data.

Table 1: Detailed Description of News Data.

Category	Time Span	Number	Min Length	Avg Length	Max Length
Train	2020/10-2022/10	169482	23	715	1407
Validation	2022/11-2023/07	58,423	35	700	1366
Test	2023/08-2024/01	34,212	27	725	1424
Total	2020/06-2024/01	252,117	23	723	1424

4.2. Evaluation Metrics

In this study, we chose precision, recall, F1 score, directional accuracy (DA), and Matthews correlation coefficient (MCC) as quantitative evaluation metrics for the stock prediction experiments. Precision, recall, and F1 score are commonly used metrics in stock prediction tasks, while DA calculates the proportion of correctly predicted instances in the dataset. MCC is a correlation coefficient that describes the relationship between actual and predicted classifications. It takes into account true positives, false positives, true negatives, and false negatives to mitigate biases caused by data imbalance.

The formulas for computing each metric are as follows:

$$Precision = \frac{t_p}{t_p + f_p} \tag{11}$$

$$Recall = \frac{t_p}{t_p + f_n} \tag{12}$$

$$F_1 = 2 \times \frac{Recall \times Precision}{Recall + Precision} \tag{13}$$

$$DA = \frac{n_{correct}}{n_{total}} \tag{14}$$

$$MCC = \frac{t_n \times t_p - f_n \times f_p}{\sqrt{(t_p + f_n)(t_p + f_p)(t_n + f_n)(t_n + f_p)}} \tag{15}$$

Where $n_{correct}$ represents the number of correctly predicted samples, n_{total} represents the total number of days in the test set, $t_p(t_n)$ is the number of correctly predicted positive (negative) labels, and $f_p(f_n)$ is the number of incorrectly predicted positive (negative) labels. To validate the financial benefits of the model, we designed an event-driven market simulation strategy to evaluate performance by backtesting on the test dataset. We also employed the Sharpe Ratio (SR) and Maximum Drawdown Rate (MDR) as metrics to assess portfolio performance. SR evaluates the ratio of investment risk to profit by calculating the excess returns generated for each unit of total portfolio risk. MDR indicates the largest peak-to-trough decline in the value of an investment or portfolio during a specific period. The formulas for calculating SR and MDR are as follows:

$$SR = \frac{E(R_P) - R_f}{\sigma_p} \tag{16}$$

$$MDR = max(\frac{Peak - Through}{Peak}) \tag{17}$$

Where $E(R_P)$ represents the expected annualized return rate of the portfolio, Rf is the annualized risk-free return, and σ_p is the standard deviation of the portfolio's annualized return.

4.3. Baseline Models

To validate the effectiveness of the EDHAN model, we compare it with the following state-of-the-art (SOTA) models as baselines:

USEP(Ding et al. 2014): Represents news events as dense vectors and uses a neural tensor network to train events.

EDSP(Ding et al. 2015): Models news events using structured event quadruples and trains events using a dual-hidden-layer network.

KDEP(Ding et al. 2016): Enhances event representation using additional information from knowledge graphs.

HATS(Liu et al. 2020): Uses a multi-layer hierarchical attention network to measure the significance of information from various news and social media sources.

HIST(Xu et al. 2022): Predicts the impact of historical events on price fluctuations by learning stock backgrounds.

In order to conduct an ablation analysis on all major modules of EDHAN, the following variations are also constructed:

- EDHAN-N: Resets the edge weights of data links to 0 to assess the impact of stock background information on model predictions.
- EDHAN-E: Resets the edge weights of stock entity links to 0 to evaluate the interactive impact of news events on stock entities.

Table 2 displays the main parameter comparisons between the EDHAN model and the baseline models.

Table 2: Comparison of Model and Baseline Accuracy.

Parameter	USEP	EDSP	KDEP	HATS	HIST	EDHAN
Batch size	128	32	128	64	64	64
Hidden layers	300	200	500	300	100-25-1	32
Characteristic dim	512	200	512	512	512	300,512
Activation	Softmax	Softmax	ReLU	ReLU	ReLU	Sigmoid
Optimizer	Adam	RMSProp	Adam	Adam	Adam	Adam
Learning rate	0.001	0.001	0.025	0.001	0.001	0.001
Epochs	500	500	5	300	300	500
Early stopping	/	10	/	/	50	10
Loss function	CE	CE	MSE	CE	MSE	BCE

4.4. Experimental Results Analysis

4.4.1 Accuracy Analysis

To compare the predictive ability of the models, we present the stock prediction results of baseline models on the overall dataset. The results of each method running 12 times are shown in Table 3. From Table 3, it’s clear that HATS has a DA of 0.734, which is the best among the baselines, while HIST has an MCC of 0.0884, which is the best among the baselines. obviously, the proposed EDHAN model outperforms all the aforementioned baselines, demonstrating the superiority of the EDHAN model compared to other baseline models.

Table 3: Comparison of Model and Baseline Accuracy.

Models	Precision	Recall	F1	DA	MCC
USEP	0.514	0.615	0.560	0.714	0.0603
EDSP	0.517	0.614	0.561	0.724	0.0658
KDEP	0.541	0.624	0.580	0.721	0.0795
HATS	0.558	0.649	0.600	0.734	0.0804
HIST	0.549	0.697	0.614	0.707	0.0884
EDHAN-N	0.505	0.683	0.581	0.716	0.0601
EDHAN-E	0.512	0.695	0.590	0.701	0.0652
EDHAN	**0.593**	**0.717**	**0.649**	**0.766**	**0.0983**

Additionally, we found that the USEP models performed the worst. The USEP model uses Open Information Extraction (Open IE) technology to predict event-driven stock market trend by extracting structured events from collected news. Then, a deep convolutional neural network is applied to simulate the effect of events on stock price fluctuations. However, these two approaches only model single news event texts without incorporating other stock market data for model learning. This highlights the importance of constructing stock background information using multi-source data for event-driven analysis. In contrast, the KDEP model incorporates extra information from a knowledge graph to improve event representation, leading to better predictive performance.

On the other hand, the HATS and HIST models incorporate graph computing methods, considering the impact of event information on related stocks while learning news events, making them the best-performing baseline models. Furthermore, in the ablation experiments, we found that constructing stock background information is conducive to model learning for analyzing financial events, and that the same financial event exhibits interdependence among related stocks, thus validating our hypothesis. Overall, our approach complements traditional stock prediction models.

We also conduct an ablation study to investigate the contribution of different types of edge information in our heterogeneous graph. Specifically, EDHAN-N disables the stock background information by removing data-based edges, while EDHAN-E removes event-driven entity-level interactions. EDHAN-N shows a noticeable decline, indicating that background data such as technical and fundamental indicators play a significant role in capturing stock-specific patterns. EDHAN-E performs slightly better than EDHAN-N but still underperforms the full model. This suggests that while entity-level event interactions are valuable, they are somewhat more complementary compared to the background signals. The full model (EDHAN) consistently achieves the best results, demonstrating that both background stock information and event-based entity relationships are essential and mutually reinforcing components of the proposed architecture.

4.4.2 Effectiveness Analysis

To further measure the real-time effectiveness of the EDHAN model, we conducted market investment analysis using test data from November 2022 to July 2023. We designed a news event monitor to simulate capturing news events and devised an event-driven investment trading strategy. Buy actions were executed on the top ten ranked stocks predicted by the model, with all stocks set to have a maximum holding period of one day, i.e., buying on the day, selling on the next day.

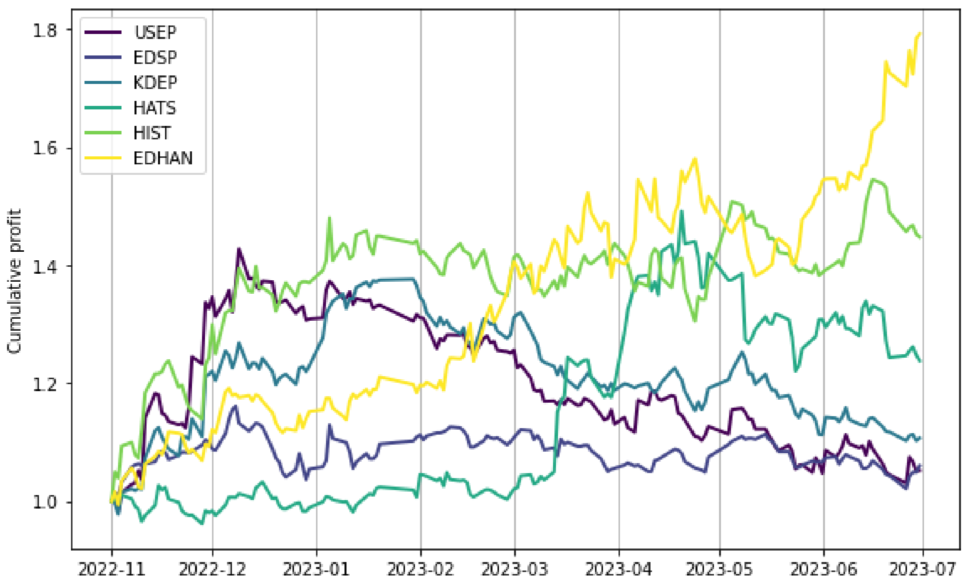

Fig 4: Cumulative Returns Comparison between Model and Baseline.

Fig. 4 depicts the cumulative returns comparison between the EDHAN model and other models. From the data in the graph, we can observe that the EDHAN model achieved the highest cumulative returns during the nine months of simulated trading, with a final profit of 78.39%, significantly outperforming the other baseline models over the same period. It is noteworthy that considering the interrelated effects of financial events on related stocks can greatly increase the cumulative returns of the model, as demonstrated by the HATS and HIST models. This is consistent with the reality where positive news often drives up the stock prices of many related companies, leading to higher excess returns from short-term trading with related stocks as the investment portfolio. Compared to other baselines, the HATS and HIST models obtained returns of 46.24% and 23.74% respectively, further confirming the effectiveness of our approach.

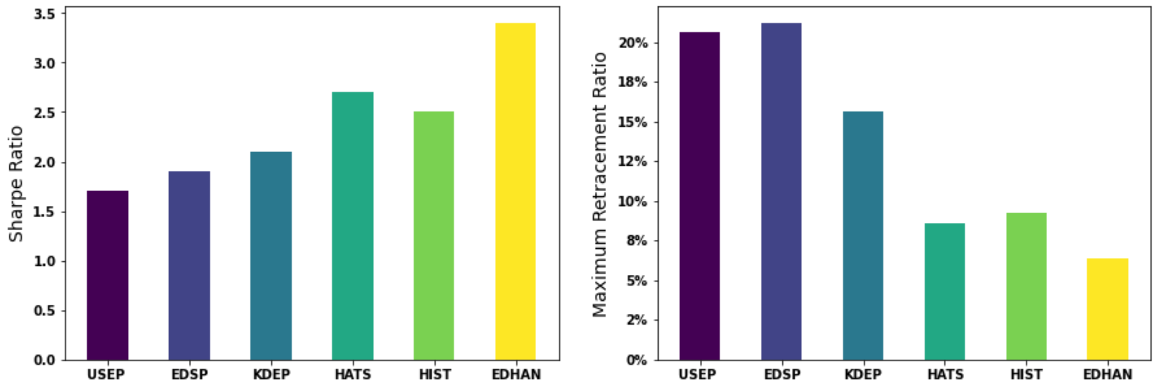

Fig 5: Compares the Sharpe Ratio (SR) and Maximum Drawdown Rate (MRR).

Fig. 5 presents a comparison of the Sharpe Ratio (SR) and Maximum Drawdown Rate (MDR) for each model. In the nine-month investment simulation experiment, the USEP models exhibit lower Sharpe ratios and higher maximum drawdown rates, indicating inaccurate predictions of financial events. The HIST model demonstrates a higher Sharpe ratio but also a higher maximum drawdown rate, suggesting insensitivity to favorable news and an inability to seize good investment opportunities. In contrast, our proposed EDHAN model shows a Sharpe ratio of 3.4 and a maximum drawdown rate of 6.4%, outperforming all baseline models. This implies that the E DHAN model achieves higher returns while having lower investment risk.

5. Conclusion

This paper presents EDHAN, a novel multi-source event-driven stock prediction model based on a heterogeneous graph attention network. By integrating technical indicators, fundamental data, and financial news, EDHAN effectively captures both the heterogeneous impact of events on different stocks and the inter-stock transmission of event effects. Extensive experiments on real-world data demonstrate its superior predictive performance and practical investment value.

Despite its strong empirical results, EDHAN has certain limitations. It currently relies on fixed temporal windows for event influence and lacks integration of broader contextual factors such as macroeconomic trends. Future work will focus on modeling dynamic temporal dependencies, incorporating richer contextual signals, and extending the framework to other financial scenarios such as policy forecasting and sector-level trend analysis.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (No.61902349) and the Zhejiang Higher Education Institutions (No.20230N082)

References

Baker, S. R., Bloom, N., & Davis, S. J. (2016). *Measuring economic policy uncertainty*. The Quarterly Journal of Economics, 131(4), 1593–1636.

Bekaert, G., Hoerova, M., & Duca, M. L. (2013). *Risk, uncertainty and monetary policy*. Journal of Monetary Economics, 60(7), 771–788.

Ceskoutsé, R. F. T., Bomgni, A. B., Zanfack, D. R. G., Agany, D. D., Thomas, B. B., & Zohim, E. G. (2024). HeteroKGRep: Heterogeneous Knowledge Graph based Drug Repositioning. *Knowledge-Based Systems*, *305*, 112638.

Chen, J., Gong, Z., Li, Y., Zhang, H., Yu, H., Zhu, J., ... & Wu, K. (2022). Meta-path based neighbors for behavioral target generalization in sequential recommendation. *IEEE Transactions on Network Science and Engineering*, *9*(3), 1658-1667.

Ding, X., Zhang, Y., Liu, T., & Duan, J. (2014, October). Using structured events to predict stock price movement: An empirical investigation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)* (pp. 1415-1425).

Ding, X., Zhang, Y., Liu, T., & Duan, J. (2015, July). Deep learning for event-driven stock prediction. In *Ijcai* (Vol. 15, pp. 2327-2333).

Ding, X., Zhang, Y., Liu, T., & Duan, J. (2016, December). Knowledge-driven event embedding for stock prediction. In *Proceedings of coling 2016, the 26th international conference on computational linguistics: Technical papers* (pp. 2133-2142).

Etzioni, O., Banko, M., Soderland, S., & Weld, D. S. (2008). Open information extraction from the web. *Communications of the ACM*, *51*(12), 68-74.

Fama, E. F. (1995). Random walks in stock market prices. *Financial analysts journal*, *51*(1), 75-80.

Jiang, Y., Li, X., Wang, Y., & Zhou, J. (2022). *Multi-task learning for event-driven stock prediction with pre-trained language models*. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1234–1245.

Jintao, L., Hongfei, L., & Liang, Y. (2020). Multi-element hierarchical attention capsule network for stock prediction [J]. *IEEE Access*, *8*, 143114-143123.

Liu, R., Liu, H., Huang, H., Song, B., & Wu, Q. (2024). Multimodal multiscale dynamic graph convolution networks for stock price prediction. *Pattern Recognition*, *149*, 110211.

Lu, M., & Xu, X. (2024). TRNN: An efficient time-series recurrent neural network for stock price prediction. *Information Sciences*, *657*, 119951.

Ma, Y., Mao, R., Lin, Q., Wu, P., & Cambria, E. (2023). Multi-source aggregated classification for stock price movement prediction. *Information Fusion*, *91*, 515-528.

Naik, N., & Mohan, B. R. (2019, February). Optimal feature selection of technical indicator and stock prediction using machine learning technique. In *International Conference on Emerging Technologies in Computer Engineering* (pp. 261-268). Singapore: Springer Singapore.

Reboredo, J. C., Rivera-Castro, M. A., & Zebende, G. F. (2017). *Oil and US dollar exchange rate dependence: A detrended cross-correlation approach*. Energy Economics, 62, 282–292.

Shi, Z., Zhang, W., Huang, J., & Ding, W. (2025). JLR-GCN: Joint label-aware and relation-aware graph convolution neural network for heterogeneous graph representations. *Information Sciences*, *706*, 122011.

Sun, Z., Yu, J., Zhang, Y., & Xu, H. (2022). *Knowledge-aware graph neural networks for event-driven stock prediction*. IEEE Transactions on Knowledge and Data Engineering.

Tan, J., Li, Q., Wang, J., & Chen, J. (2022). FinHGNN: A conditional heterogeneous graph learning to address relational attributes for stock predictions. *Information Sciences*, *618*, 317-335.

Tetlock, P. C. (2007). *Giving content to investor sentiment: The role of media in the stock market*. The Journal of Finance, 62(3), 1139–1168.

Wang, T., Pan, Z., Hu, G., Xu, K., & Zhang, Y. (2023). Self-supervised heterogeneous graph learning with iterative similarity distillation. *Knowledge-Based Systems*, *276*, 110779.

Wang, X., Ji, H., Shi, C., Wang, B., Ye, Y., Cui, P., & Yu, P. S. (2019, May). Heterogeneous graph attention network. In *The world wide web conference* (pp. 2022-2032).

Wu, Q., Liu, M., & Gao, W. (2023). *Event-aware heterogeneous graph learning for stock prediction*. Knowledge-Based Systems, 256, 109998.

Xu, W., Liu, W., Wang, L., Xia, Y., Bian, J., Yin, J., & Liu, T. Y. (2021). Hist: A graph-based framework for stock trend forecasting via mining concept-oriented shared information. *arXiv preprint arXiv:2110.13716*.

Yang, A., Li, M., Ding, Y., Fang, L., Cai, Y., & He, Y. (2024). Graphmamba: An efficient graph structure learning vision mamba for hyperspectral image classification. *IEEE Transactions on Geoscience and Remote Sensing*.

Yang, F., Zhu, C., & Lin, H. (2023). *Causal event modeling for financial news-driven stock prediction*. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(6), 5842–5850.

Yang, H., Zhu, Y., & Huang, Q. (2018). A multi-indicator feature selection for CNN-driven stock index prediction. In *Neural Information Processing: 25th International Conference, ICONIP 2018, Siem Reap, Cambodia, December 13–16, 2018, Proceedings, Part V 25* (pp. 35-46). Springer International Publishing.

Yang, X., Yan, M., Pan, S., Ye, X., & Fan, D. (2023, June). Simple and efficient heterogeneous graph neural network. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 37, No. 9, pp. 10816-10824).

Zhao, Z., Zhu, Z., Liu, Y., Guo, J., & Yang, K. (2024). Heterogeneous graph contrastive learning with augmentation graph. *IEEE Transactions on Artificial Intelligence*.

Zou, J., Cao, H., Liu, L., Lin, Y., Abbasnejad, E., & Shi, J. Q. (2022). Astock: A new dataset and automated stock trading based on stock-specific news analyzing model. *arXiv preprint arXiv:2206.06606*.

Appendix

1. Technical And Fundamental Features Used in The Model

Feature Type	Indicators Metrics	Description
Technical Features (x_i)	Closing Price	Daily closing price of the stock
	Daily Return Rate	$\frac{P_t - P_{t-1}}{P_{t-1}}$, where P_t is the closing price at time t
	Moving Average (5-day, 10-day)	Short-term trend indicators based on recent closing prices
	Volatility (5-day Std. Dev.)	Measures recent price fluctuations
	RSI	Relative Strength Index; detects overbought/oversold conditions
	MACD	Moving Average Convergence Divergence; momentum indicator
	Turnover Rate	Trading volume normalized by float shares
Fundamental Features (z_i)	P/E Ratio	Price-to-Earnings ratio; valuation measure
	EPS	Earnings Per Share; company profitability
	ROE	Return on Equity; financial efficiency
	Debt-to-Equity Ratio	Leverage ratio indicating financial risk
	Operating Revenue Growth	Year-over-year revenue growth
	Net Profit Margin	Net income divided by revenue
	Market Capitalization	Total market value of the company's outstanding shares