Multi-Source Aggregated Classification for Stock Price Movement Prediction

YU MA[1], RUI MAO[2], QIKA LIN[3], PENG WU[4], and ERIK CAMBRIA[2]
[1]School of Economics and Management, Nanjing University of Science and Technology, 200 Xiaolingwei, Nanjing 210094, Jiangsu, China
[2]School of Computer Science and Engineering, Nanyang Technological University, 50 Nanyang Avenue, 639798, Singapore
[3]School of Computer Science and Technology, Xi'an Jiaotong University, Xi'an 710049, Shaanxi, China
[4]School of Intelligent Manufacturing, Nanjing University of Science and Technology, 200 Xiaolingwei, Nanjing 210094, Jiangsu, China
Corresponding author: Yu Ma.

ABSTRACT Stock price movement prediction remains a challenging research task. Existing studies mainly rely on numerical features of the target stock and news sentiment; however, semantics-based sentiment analysis cannot accurately capture real market sentiment. In addition, using only information about the target company is insufficient for explaining stock price fluctuations, because the price of a target stock may also be affected by related companies. To address these limitations, this paper proposes a Multi-source Aggregated Classification (MAC) method for stock price movement prediction. The proposed method jointly exploits numerical features of the target stock, market-driven news sentiment, and news sentiment from related stocks. To extract real market sentiment from news more effectively, we design a pretraining task for an embedding generator supervised by actual stock price movements, enabling the embeddings produced by the pretrained sentiment classifier to encode market sentiment information in vector space. Furthermore, MAC introduces a graph convolutional network to model the influence of news about related companies on the target stock and then predicts the stock price movement on the next trading day. Experimental results show that MAC outperforms the baselines considered in this study in stock price movement prediction, Sharpe ratio, and backtesting return.

INDEX TERMS stock price movement prediction, financial news, graph convolutional network, BiLSTM, market-driven sentiment, multi-source fusion

I. INTRODUCTION

In recent years, the rapid development of artificial intelligence and machine learning has had a profound impact on the financial domain, especially on quantitative stock prediction research [1-3]. Developing effective stock price movement prediction methods can help investors identify investment opportunities and help securities institutions reduce the risks associated with margin trading and securities lending. Because the stock market is highly complex and stock prices are usually affected by multiple factors [4,5], how to aggregate and exploit multi-source information to improve prediction performance has remained a central concern in both academia and industry.

For stock price movement prediction, prior studies have widely used quantitative indicators, including historical trading data and technical indicators [2,6], and have demonstrated their effectiveness for prediction tasks. In addition to quantitative information, event-driven trading is also an important investment paradigm because stock price movements are influenced by the release, diffusion, and absorption of information [7-10]. In this context, financial news has become an important source of information for investors. It affects investor sentiment and trading behavior and, in turn, influences stock price changes [11,12]. With the development of natural language processing (NLP), an increasing number of studies have combined financial news sentiment analysis with quantitative indicators for stock price movement prediction [13-15]. Among them, predefined sentiment lexicons are widely used to analyze the sentiment of financial news, and the resulting lexical sentiment scores are directly used as sentiment features [14-18]. However, lexicon-based methods are highly dependent on the quality and coverage of the lexicon [19-22], and they have difficulty adapting to dynamic sentiment semantics under different market conditions such as bull and bear markets.

In recent years, deep learning models have also been introduced into financial news sentiment analysis to further improve stock prediction performance [23-25]. These methods usually rely on labeled data to learn sentiment polarities such as positive, neutral, and negative, and then fit stock price movements accordingly. However, such methods still suffer from two limitations. First, constructing large-scale domain-specific sentiment analysis datasets is labor-intensive and time-consuming. Second, sentiment polarities derived from semantic interpretation and manual annotation may not faithfully reflect real market sentiment [26]. This is because human annotators are not the actual investors whose trading behavior moves market prices in real time, and therefore there remains an evident gap between semantic judgments and real market responses.

Most existing event-driven stock prediction studies focus on a one-to-one scenario, implicitly assuming that a company's stock price is affected only by news about that company, while paying less attention to the spillover effects of news about related companies [14-16]. However, with the development of the market economy, listed companies have formed complex relation networks through supply chains [27], shareholding relations [28], and industry competition [29]. Research on crisis spillover effects shows that the negative impact of a crisis event in one company may not only affect the company itself, but may also spread to other related firms [30]. Therefore, incorporating news about related companies into the prediction of a target company's stock price movement is necessary.

To address the above issues, this paper proposes the MAC model for stock price movement prediction. From the perspective of investment decision-making, stock market investment mainly includes portfolio management and stock trading [31]. Portfolio management focuses on the allocation and management of a group of stocks in long-term investment scenarios, whereas stock trading pays more attention to short-term strategies for individual stocks. This study focuses on the stock trading task. By predicting the price movement of a specific stock on the next trading day, the proposed MAC model provides decision support for investors. If the model predicts an upward signal for the next trading day, investors may buy or continue to hold the stock on the current day. If the model predicts a downward signal, investors may choose not to buy or not to continue holding the stock.

The MAC model integrates three sources of features: trading data and technical indicators of the target stock, news about the target stock, and news about stocks related to the target stock. We select 31 robust quantitative indicators to characterize stock price movement patterns. For news information, we construct a pretrained embedding generator based on Chinese RoBERTa [32,33]. Its training objective is to fit the actual stock price movement of a company on the next trading day using company news, so that the resulting news embeddings can better reflect market sentiment, where upward movement corresponds to positive sentiment and downward movement corresponds to negative sentiment. On this basis, a graph convolutional network (GCN) [34] is further adopted to model the connections between the target company and related companies. Finally, the multi-source features are concatenated and fed into a Bidirectional Long Short-Term Memory (BiLSTM) network [35] to predict the stock price movement on the next trading day.

We evaluate the effectiveness of MAC on six target stocks from six different industries in the Chinese stock market. Experimental results show that, on the binary stock price movement classification task, MAC improves average accuracy by 2.38% and Matthews correlation coefficient by 4.62% over eLSTM [17], the strongest baseline in our comparative experiments. In financial evaluation, the average Sharpe ratio of MAC is 0.68 higher than that of eLSTM, and backtesting also yields stronger trading returns. In addition, we conduct systematic ablation studies to verify the effectiveness of different feature sources and different pretraining strategies, namely market-driven sentiment pretraining versus semantics-based sentiment pretraining.

The main contributions of this paper are summarized as follows:

- We propose a multi-source fusion model for stock price movement prediction that incorporates both target-stock information and news about related companies.
- We design a market-driven sentiment pretraining task to learn vector-space representations of news.
- Experimental results show that, under our experimental settings, the proposed model outperforms the compared baselines, and that market-driven sentiment pretraining also outperforms semantics-based sentiment pretraining.

The rest of this paper is organized as follows. Section 2 reviews related work. Section 3 introduces the proposed MAC method. Section 4 presents the experimental setup. Section 5 reports experimental results, ablation studies, hyperparameter analysis, and fine-grained analysis. Section 6 concludes the paper and outlines future work.

II. RELATED WORK

A. INFORMATION AND STOCK PRICE MOVEMENTS

Stock price movements are essentially the result of the combined effects of the release, diffusion, and absorption of multi-source information [7,37]. Among different types of information, major news events are an important source of return volatility [9]. Jeon et al. [9] pointed out that stock return volatility is significantly related to both the amount and content of news. Shi and Ho [38] found that negative news events can lead to abnormal short-term stock price volatility. Strauss et al. [39] argued that excessive media attention and negative news can have a significantly negative impact on stock returns. In general, news events are an important external factor driving stock price movements.

With the development of the market economy, listed companies have formed extensive and complex relation networks through supply chains [27], shareholding chains [28], and industry competition [29]. In this context, the stock price movement of a company may be influenced not only by its own news, but also indirectly by news about related companies. Research in marketing and crisis management has shown that spillover effects commonly exist among related companies, that is, information about company A can affect the public perception and judgment of company B [40], and the magnitude of the spillover effect depends on the degree of association between the two companies [40]. Jacobs and Singhal [41] found that negative public opinion about a listed company exhibits a strong vertical spillover effect along the supply chain, which may cause the stock prices of its suppliers and customers to decline simultaneously. Moreover, when a parent brand encounters adverse events, sub-brands also become vulnerable [42], and vice versa [40]. Existing studies have further shown that brand crises in listed companies may spill over to competitors [29,41,43], thereby affecting the entire industry. Therefore, financial news about related companies connected through supply-chain, shareholding, and competitive relations is an important information source for predicting the stock price movement of a target company.

In summary, news-driven stock prediction should consider not only news about the target company itself, but also news about related companies.

B. STOCK PRICE MOVEMENT PREDICTION

Stock price movement prediction has long been a hot research topic in both academia and industry [1,2]. In relevant studies, quantitative indicators are widely used as basic features, including daily price data and technical indicators such as moving averages (MA), Williams %R (WR), and the relative strength index (RSI) [1,6,28]. These features can characterize, to some extent, the shape of stock time series and trading behavior patterns.

With the development of NLP, an increasing number of studies have combined financial news with quantitative indicators based on behavioral finance theory to improve stock prediction performance [13,44]. Some of these studies use sentiment lexicons to extract features from financial news and then use them for stock prediction [14-17,45]. For example, Chen et al. [15] fed quantitative indicators and news features extracted from a sentiment lexicon into the RNN-boost model, thereby improving stock prediction in the Chinese market. Chen et al. [16] incorporated daily prices and news sentiment features into an LSTM [46] model and further improved stock prediction performance. Li et al. [14] combined quantitative indicators with news sentiment extracted from a sentiment lexicon in an LSTM-based model for predicting the Hong Kong stock market, and pointed out that domain-specific financial sentiment lexicons outperform approaches based only on technical indicators or general-purpose sentiment lexicons.

In addition to lexicon-based methods, machine-learning approaches to extracting financial news features have also been widely used in stock prediction [1,47]. These studies typically rely on manually labeled positive and negative sentences to train financial news sentiment classifiers [23-25]. For example, de Oliveira Carosia et al. [23] trained an ANN classifier on Brazilian Portuguese financial news to extract news features, where the news corpus was manually annotated by experts according to semantic sentiment. Sousa et al. [24] manually labeled 582 stock news articles as positive, neutral, or negative, and fine-tuned BERT [48] for financial news sentiment analysis.

Although the above studies have made encouraging progress, they still have several limitations. On the one hand, lexicon-based methods usually construct sentiment features using positive and negative word frequencies or lexical sentiment scores, but they are highly dependent on lexicon quality and coverage and have difficulty modeling dependency relations in the overall context [21]. In addition, the construction of sentiment lexicons depends heavily on manual curation, which makes it difficult to adapt in a timely manner to sentiment changes under different market conditions, such as bull and bear markets. On the other hand, machine-learning methods usually perform supervised learning on manually annotated sentiment analysis datasets to obtain sentiment features. However, the labels in these datasets are typically derived from semantic interpretation of the text, and the annotators are not real market participants whose actions directly affect stock prices. Therefore, a gap still exists between semantics-based sentiment features and real market sentiment. More importantly, although both empirical studies and theoretical analyses have shown that news about related companies affects the stock prices of target companies, neither lexicon-based methods nor machine-learning methods generally incorporate news about related companies into the modeling process.

III. METHOD

This section introduces the proposed multi-source information fusion model, whose input consists of quantitative indicators of the target stock, news about the target stock, and news about related stocks. Following previous studies [14,17], we define the task as predicting the direction of the target stock price movement on the next trading day, namely upward or downward. If $CP_{t+1} - CP_t \geqslant 0$, the label y_{t+1} is set to 1; otherwise, it is set to 0. Here, CP_{t+1} denotes the closing price of the target stock on trading day $t+1$, and CP_t denotes the closing price on day t. We do not introduce an additional label for the case $CP_{t+1} - CP_t = 0$, because the situation in which the closing prices on two consecutive trading days are exactly the same is relatively rare.

As shown in Fig. 1, the proposed MAC model consists of four technical modules. First (Section 3.1), it extracts quantitative indicators of the target stock and their subsequences within a time window T. Here, the target stock is the stock whose price movement on the next trading day needs to be predicted, T denotes the number of trading days before the prediction date, and Q_{tar} denotes the quantitative indicators of the target stock within the T trading days. Second (Section 3.2), it constructs target-stock news features N_{tar} over the T trading days. This feature is generated by an embedding generator, which is essentially a pretrained market-driven sentiment classifier based on Chinese RoBERTa [33]. This classifier extracts an embedding for each news headline and averages the embeddings of all headlines published on the same trading day, thereby obtaining the representation of the target stock within the time window T. Third (Section 3.3), it identifies important related stocks for the target stock according to cosine similarity between stock nodes. The news features of related stocks are also generated by the pretrained classifier and averaged by day to obtain the news representations of important related stocks within the time window T, denoted by N_{rel}. GCN is then used to model the relations among stocks and aggregate news representations from related stocks, yielding the related-stock news representation of the target stock, R_{tar}. Fourth (Section 3.4), the multi-source information Q_{tar}, N_{tar}, and R_{tar} is concatenated and fed into a BiLSTM-based classifier to predict the stock price movement on the next trading day, y_{t+1}. The following subsections describe these modules in detail, and Table 1 summarizes the symbols used in this section.

TABLE 1. Notation used in this section.

Symbol	Description
CP_t	Closing price of the target stock on day t.
T	Number of trading days before the stock movement prediction day.
K	Top-K stocks most related to the target stock.
$q_{tar,t}$	$q_{tar,t} \in \mathbb{R}^{1\times 31}$, quantitative indicators of the target stock on day t.
Q_{tar}	$Q_{tar} \in \mathbb{R}^{T\times 31}$, quantitative indicators of the target stock within time window T.
e	Market-driven sentiment embedding of a single news headline.
n	Daily news feature.
$MSC(\cdot)$	News embedding generator based on the pretrained market-driven sentiment classifier.
$n_{tar,t}$	$n_{tar,t} \in \mathbb{R}^{1\times 768}$, news feature of the target stock on day t.
N_{tar}	$N_{tar} \in \mathbb{R}^{T\times 200}$, news features of the target stock within time window T.
$G = (V, E)$	Stock relation graph, where V is the set of nodes (stocks) and E is the set of edges (relations).
$N_{rel,t}$	$N_{rel,t} \in \mathbb{R}^{K\times 768}$, news features of the K related stocks on day t.
N_{rel}	$N_{rel} \in \mathbb{R}^{(T\times K)\times 768}$, news features of the K related stocks within time window T.
A	$A \in \mathbb{R}^{K\times K}$, adjacency matrix.
M	$M \in \mathbb{R}^{(T\times K)\times(T\times K)}$, adjacency matrix over T time steps.
H	$H \in \mathbb{R}^{(T\times K)\times 200}$, GCN hidden state representing related-stock features.
R_{tar}	$R_{tar} \in \mathbb{R}^{T\times 200}$, related-stock news features of the target stock within time window T.

A. EXTRACTION OF QUANTITATIVE INDICATORS FOR THE TARGET STOCK

Quantitative indicators are among the most common numerical features in stock prediction and are typically composed of historical trading data and technical indicators. Following previous studies [14,17,45], we select 31 commonly used quantitative indicators, as listed in Table 2. The technical indicators are computed using the formulas in the Appendix, and the historical quantitative data of the target stock are obtained from the RESSET platform. Because different indicators have different value ranges, Min-Max normalization [45] is applied to all quantitative indicators. After processing, the quantitative indicators of the target stock (tar) on trading day t can be expressed as:

$$q_{tar,t} = [OP_{tar,t}, CP_{tar,t}, \ldots, LowerBand_{tar,t}], \tag{1}$$

where $q_{tar,t} \in \mathbb{R}^{1\times 31}$ is a 31-dimensional vector. We collect quantitative indicators of the target stock over n trading days, where n denotes the time span covered by the training and testing sets. For each training step that predicts the stock movement on day $t+1$ (y_{t+1}), the quantitative indicators from the T trading days prior to day $t+1$ are used to represent temporal information, where $T \le n$. We do not directly use the entire historical series starting from day 1, because overly distant data may introduce additional noise into the prediction of y_{t+1}. The effect of different values of T is analyzed in Section 5.4. Therefore, the quantitative indicators within time window T can be expressed as:

$$Q_{tar} = [q_{tar,t-T+1}, \ldots, q_{tar,t-1}, q_{tar,t}], \tag{2}$$

where $Q_{tar} \in \mathbb{R}^{T\times 31}$ is the matrix of quantitative indicator features of the target stock.

B. EXTRACTION OF TARGET-STOCK NEWS FEATURES

This section first introduces the pretrained market-driven sentiment classifier (MSC) for generating news embeddings (Section 3.2.1) and then describes how target-stock news features are constructed based on MSC (Section 3.2.2).

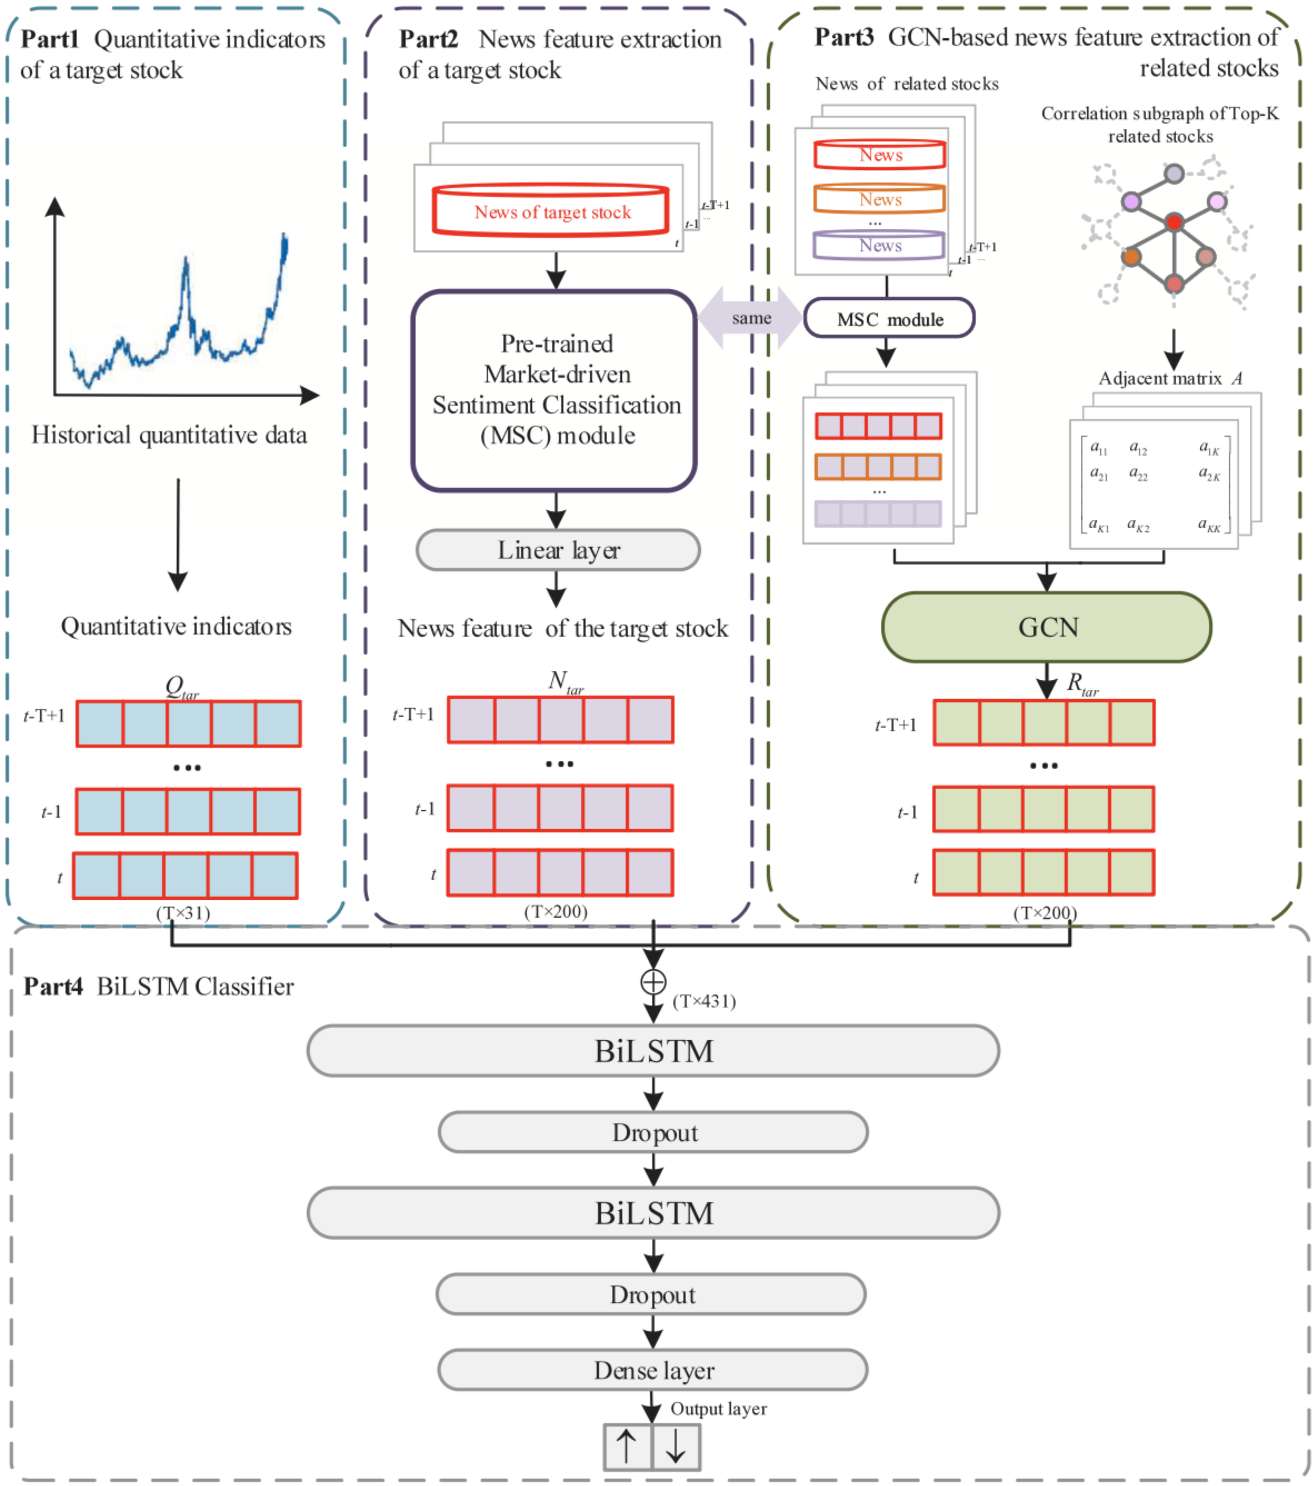

FIGURE 1. Fig. 1: Overview of the MAC framework

TABLE 2. Quantitative indicators used for the target stock.

Indicator	Indicator
Open/Close Price (OP/CP)	Moving Averages (MA10/MA20/MA30)
High/Low Price (HP/LP)	Weighted Moving Average (WMA)
Volume (Vol)	Exponential Moving Average (EMA)
Turnover Rate (TurR)	Relative Strength Index (RSI6/RSI12/RSI24)
Daily Amplitude (DA)	Money Flow Index (MFI)
Daily Return (DR)	Stochastic Oscillator KDJ (Slowk/Slowd)
Price-to-Earnings Ratio (P/E)	True Range (TR)
Price-to-Book Ratio (P/B)	Average True Range (ATR)
Price-to-Sales Ratio (P/S)	Williams %R (WR)
MACD (DIFF/DEA/MACD)	Bollinger Bands (UpperBand/MiddleBand/LowerBand)

Pretrained market-driven sentiment classifier. This subsection aims to pretrain a market-driven sentiment classifier for generating effective news embedding features. The motivation is that there is a gap between semantic sentiment and market sentiment, because the semantic content of financial news does not directly reflect investment behavior. For example, a decrease in the debt ratio and a decrease in investment amount may have very different implications for the same company, although semantically both statements indicate a reduction in a financial indicator. Specifically, investors often regard a decrease in the debt ratio as positive, because a lower debt ratio implies lower operational risk. By contrast, a reduction in investment, such as a decrease in marketing or R&D spending, may be seen as a signal that the company's development is being constrained. Therefore, we assume that stock price movement can represent investors' real sentiment toward a given piece of news. We use news headlines as input because they summarize the main content of the news. This setting is consistent with related studies, which have shown that, in stock prediction tasks, news headlines are more suitable than full news articles for accurately extracting key information [49,50]. The objective of pretraining is as follows: given the news headlines of a company on the current day, predict the stock price movement of that company on the next trading day. To this end, we first collect 24,000 Chinese financial news items corresponding to 32 stocks from January 2, 2018 to October 12, 2020, with the data provided by Hundsun. Hundsun (https://www.hundsun.com) is a leading Chinese fintech company that provides one-stop financial technology services. According to Hundsun's classification standard, the collected news can be divided into three categories: company information, business information, and financial information. Company information involves business and management changes, qualifications and honors, and external credit ratings. Business information covers operational risks, products, projects, cooperation, contracts, business activities, and violations. Financial information includes equity assets, corporate finance, financing transactions, and debt. To avoid information leakage, the pretraining dataset does not include the target companies in the downstream task or their related companies. We then label each news headline as positive (upward and unchanged) or negative (downward) according to the actual stock price movement on the next trading day after the news is published. We adopt Chinese RoBERTa [33] as the encoder because previous studies have shown that RoBERTa performs well in a variety of NLP tasks [51-56]. We prepend and append the news headline with [CLS] and [SEP], respectively, which are predefined special tokens in RoBERTa. The padded sequence is then input into the RoBERTa encoder $enc(\cdot)$ to obtain the following hidden state S:

$$S = enc(CLS, w_1, w_2, ..., w_l, SEP), \tag{3}$$

where w denotes a token in the news headline, l is the length of the input headline, and $S \in \mathbb{R}^{l\times 768}$. We take the embedding at the [CLS] position in S, denoted by e, as the embedding representation of the news headline:

$$e = S_{[CLS]}, \tag{4}$$

where $e \in \mathbb{R}^{1\times 768}$. The embedding e is then fed into a feed-forward neural network (FNN) and a softmax layer to predict the market-driven sentiment label $\hat{y}_{pre}$ during pretraining:

$$\hat{y}_{pre} = softmax(FNN(e)). \tag{5}$$

Cross-entropy loss is used in the pretraining stage:

$$\mathcal{L}_{pre} = CrossEntropy(\hat{y}_{pre}, y_{pre}), \tag{6}$$

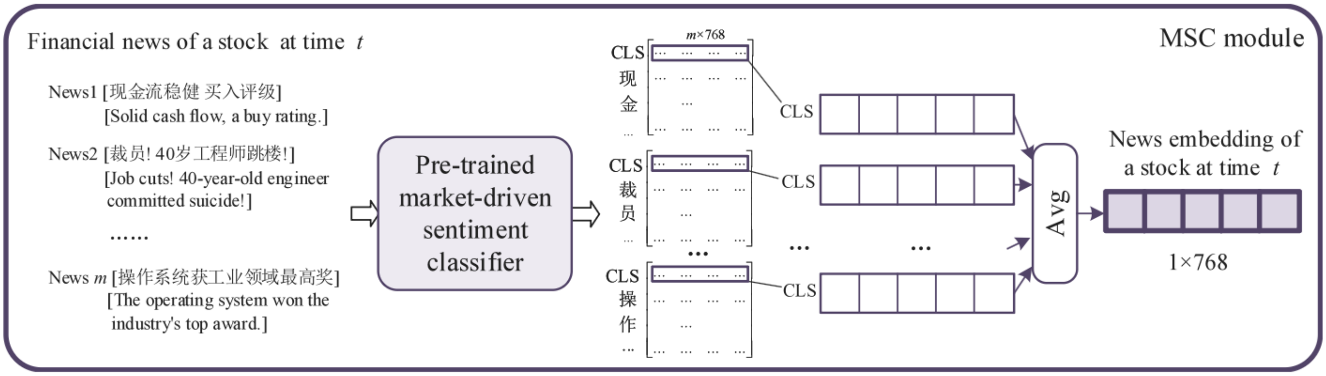

FIGURE 2. **Fig. 2: MSC module for extracting daily financial news features**

where y_{pre} is the ground-truth label in pretraining, namely positive or negative. In this way, the pretrained sentiment classifier can generate a market-driven sentiment embedding e for a given news headline. In the downstream task, a stock usually corresponds to multiple news items on the same trading day, so we average the news embeddings to obtain a daily news feature representing all headlines on that day:

$$n = \frac{1}{m} \sum_{i=1}^{m} e_i, \tag{7}$$

where m is the number of news headlines associated with the stock on that day. The process of generating news embeddings is shown in Fig. 2. For simplicity, Eqs. (3), (4), and (7) are combined. Given a set of news headlines for a stock on a trading day, the daily news feature can be expressed as:

$$n = MSC(news_1, news_2, ..., news_m), \tag{8}$$

where $MSC(\cdot)$ denotes the news embedding generator based on the pretrained market-driven sentiment classifier.

Target-stock news features. We use the pretrained MSC module, i.e., $MSC(\cdot)$ in Eq. (8), to extract the news feature of the target stock on day t, denoted by $n_{tar,t} \in \mathbb{R}^{1\times 768}$. If no financial news about the target stock is available on a given trading day, the corresponding news feature is set to a zero vector. We then construct news features over n trading days for training and testing. Consistent with the processing of quantitative indicators, for each training step we extract news features from the T trading days before the prediction day $t+1$. Meanwhile, MAC adds a linear layer, $Linear(\cdot)$, after the MSC output to reduce the dimensionality of target-stock news features. Therefore, the news features of the target company can be expressed as:

$$N_{tar} = Linear([n_{tar,t-T+1}, ..., n_{tar,t-1}, n_{tar,t}]), \tag{9}$$

where $N_{tar} \in \mathbb{R}^{T\times 200}$ is the news feature matrix of the target stock. The parameters of $Linear(\cdot)$ are updated during MAC training.

C. EXTRACTION OF RELATED-STOCK NEWS FEATURES

The extraction of related-stock news features mainly consists of three steps: related-stock discovery (Section 3.3.1), construction of related-stock news features (Section 3.3.2), and GCN-based fusion of related-stock information (Section 3.3.3).

Discovery of related stocks. The goal of related-stock discovery is to identify the stocks most related to the target stock. Selecting only the most related stocks helps avoid introducing additional noise from weakly related news and also reduces computational cost. The stock relation data are collected from Qichacha. Qichacha (https://www.qcc.com/) is a leading Chinese enterprise information service platform that mainly provides corporate credit and association information. In total, we collect relation data for 3,495 stocks in the Chinese stock market. For each stock, we gather information on its supply chain (suppliers and customers), shareholding chain (holding and held), and industry competition. Based on these data, we further identify the stocks most related to the target stock. Fig. 3 illustrates the process of related-stock discovery and the construction of the adjacency matrix for the most related stocks. First, we build a stock relation graph based on the collected relation data. The graph, denoted by $G = (V, E)$, contains 3,495 nodes V and 22,886 edges E, where each node represents a stock and each edge represents a relation between two stocks. We then apply the node2vec [57] algorithm to embed the graph and obtain vector representations for all nodes. Given the target stock and its node2vec embedding, we identify the top-K stocks most related to

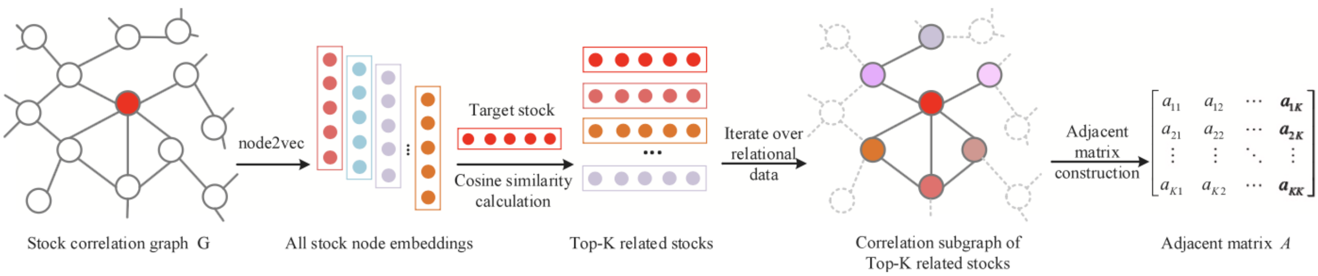

FIGURE 3. Fig. 3: Related-stock discovery and adjacency matrix construction

the target stock using cosine similarity. Among these top-K stocks, the most related one is the target stock itself because it has the highest cosine similarity. Finally, we construct an adjacency matrix $A \in \mathbb{R}^{K\times K}$ to represent the relations among these K most related stocks. K is a hyperparameter whose effect is analyzed in Section 5.4. The adjacency matrix A is defined in Eq. (10). Here, $a_{i,j}$ is an element of A; if there exists an edge between stocks s_i and s_j, then $a_{i,j} = 1$, otherwise $a_{i,j} = 0$. The row and column indices of the matrix are ordered in descending order of cosine similarity to the target stock:

$$A = \begin{bmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,K} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,K} \\ \vdots & \vdots & \ddots & \vdots \\ a_{K,1} & a_{K,2} & \cdots & a_{K,K} \end{bmatrix}. \tag{10}$$

News features of related stocks. In this step, the feature matrix is constructed by incorporating financial news from related stocks. Fig. 4 shows the construction process of the related-stock feature matrix. After collecting the financial news headlines of each of the top-K related stocks on trading day t, we use the MSC module (Section 3.2.1) to extract news features for each related stock. According to Eq. (8), the news feature of the i-th important related stock on trading day t is denoted by $n_{i,t}$. Accordingly, the news feature matrix of the K important related stocks on trading day t, denoted by $N_{rel,t} \in \mathbb{R}^{K\times 768}$, can be expressed as:

$$N_{rel,t} = [n_{1,t}, n_{2,t}, ..., n_{i,t}, ..., n_{K,t}]. \tag{11}$$

For each training step, we further retrieve the news features of important related stocks within time window T, denoted by N_{rel}:

$$N_{rel} = [N_{rel,t-T+1}, ..., N_{rel,t-1}, N_{rel,t}], \tag{12}$$

where $N_{rel} \in \mathbb{R}^{(T\times K)\times 768}$. The corresponding adjacency matrix with T time steps, denoted by M, is defined as:

$$M = \mathrm{diag}(\underbrace{A, A, ..., A}_{T}), \tag{13}$$

where $M \in \mathbb{R}^{(T\times K)\times(T\times K)}$. Here, $A \in \mathbb{R}^{K\times K}$ is the adjacency matrix located on the diagonal of M (see Eq. (10)), and all other positions in M are zero.

GCN-based fusion of related-stock information. One of the major goals of MAC is to aggregate news information from related stocks. Since the relations between the target stock and related stocks have already been modeled as a graph, we adopt the GCN model proposed by Kipf and Welling [34] to learn this graph and fuse related-stock information. The reason for choosing GCN is that it can effectively capture the structural information of nodes and edges [58,59]. The inputs of GCN include: (1) a feature matrix containing the features of all nodes, and (2) an adjacency matrix representing the relations among nodes. Through convolution, the feature of each node can be updated by aggregating the representations of its neighboring nodes according to the given adjacency matrix. Fig. 5 illustrates the GCN-based fusion process for related-stock information. Specifically, we first construct the related-stock news feature matrix N_{rel} using a rolling window and use time step T as the GCN input (see Eq. (12)). At the same time, the corresponding adjacency matrix M with T time steps is constructed (see Eq. (13)). The news feature matrix N_{rel} and adjacency matrix M are then fed into the GCN layer, and the convolution operation is performed as follows:

$$H = ReLU(\tilde{D}^{-1/2} M \tilde{D}^{-1/2} N_{rel} W), \tag{14}$$

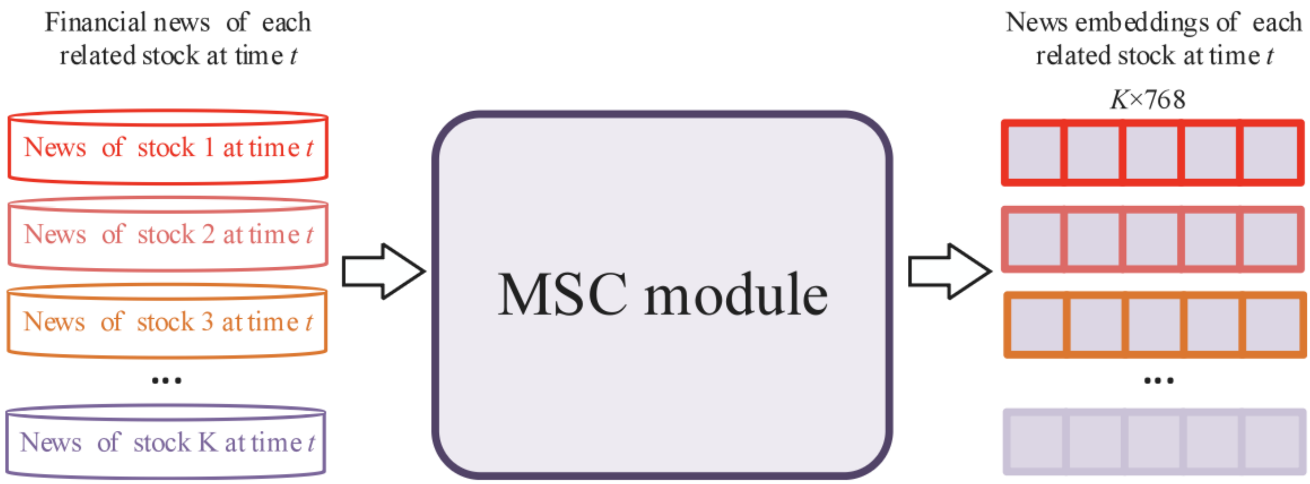

FIGURE 4. **Fig. 4: Construction of the related-stock feature matrix**

where $\tilde{D}^{-1/2}M\tilde{D}^{-1/2} \in \mathbb{R}^{(T\times K)\times(T\times K)}$ is the normalized propagation matrix, $\tilde{D} \in \mathbb{R}^{(T\times K)\times(T\times K)}$ is the diagonal degree matrix whose diagonal entries satisfy $\tilde{D}_{i,i} = \sum_j M_{i,j}$, $M \in \mathbb{R}^{(T\times K)\times(T\times K)}$ is the adjacency matrix, $N_{rel} \in \mathbb{R}^{(T\times K)\times 768}$ is the related-stock news feature matrix, $W \in \mathbb{R}^{768\times 200}$ is the trainable weight matrix of the GCN layer, $ReLU(\cdot)$ is the activation function, and $H \in \mathbb{R}^{(T\times K)\times 200}$ is the GCN output. In GCN, each node feature vector is updated by aggregating information from neighboring nodes. Thus, the updated representation of each stock node contains information from related stocks on the same trading day. We further reshape the matrix H from a two-dimensional form, $\mathbb{R}^{(T\times K)\times 200}$, into a three-dimensional form, $\mathbb{R}^{T\times K\times 200}$, expressed as:

$$H = [H_{t-T+1}, ..., H_{t-1}, H_t], \tag{15}$$

where $H_t \in \mathbb{R}^{K\times 200}$ is the matrix containing the feature vectors of the K related stocks on trading day t. H_t can be written as:

$$H_t = [n'_{1,t}, n'_{i,t}, ..., n'_{K,t}], \tag{16}$$

where $n'_{1,t} \in \mathbb{R}^{1\times 200}$ is the updated feature vector of the stock with the highest relatedness, namely the target stock itself, obtained by aggregating information from the other top-K related stocks on trading day t. Finally, the feature vector n'_1 is extracted from each trading day within window T to construct the matrix $R_{tar} \in \mathbb{R}^{T\times 200}$, which represents the sequential news features contributed by related stocks to the target stock over the past T trading days. R_{tar} is defined as:

$$R_{tar} = [n'_{1,t-T+1}, ..., n'_{1,t-1}, n'_{1,t}]. \tag{17}$$

D. STOCK PRICE MOVEMENT PREDICTION

BiLSTM is an effective sequence learning model that performs well on sequence processing tasks [60]. In this section, BiLSTM is used to capture temporal financial information for stock price movement prediction. We also test other common NLP encoders, including LSTM [46], GRU [61], Bi-GRU, and Transformer [62], and the experimental results show that BiLSTM is more robust for this task. The stock price movement prediction process based on BiLSTM is shown in Fig. 6. Specifically, given a rolling window size T, we concatenate the quantitative indicators of the target stock Q_{tar}, the target-stock news features N_{tar}, and the related-stock news features R_{tar} into X. The input to BiLSTM can thus be expressed as:

$$X = [Q_{tar} \oplus N_{tar} \oplus R_{tar}], \tag{18}$$

where $X \in \mathbb{R}^{T\times 431}$, $\oplus$ denotes concatenation, $x_{t-T+1}, ..., x_{t-1}, x_t$ are the day-wise concatenated features in X, and $x_t \in \mathbb{R}^{1\times 431}$. We then feed X into two BiLSTM layers, where the number of layers is examined in Section 5.4, to mine temporal patterns for stock price movement prediction. To reduce overfitting, dropout $\mathcal{D}(\cdot)$ is applied after each BiLSTM layer:

$$h_{t+1} = \mathcal{D}(BiLSTM_2(\mathcal{D}(BiLSTM_1(X)))), \tag{19}$$

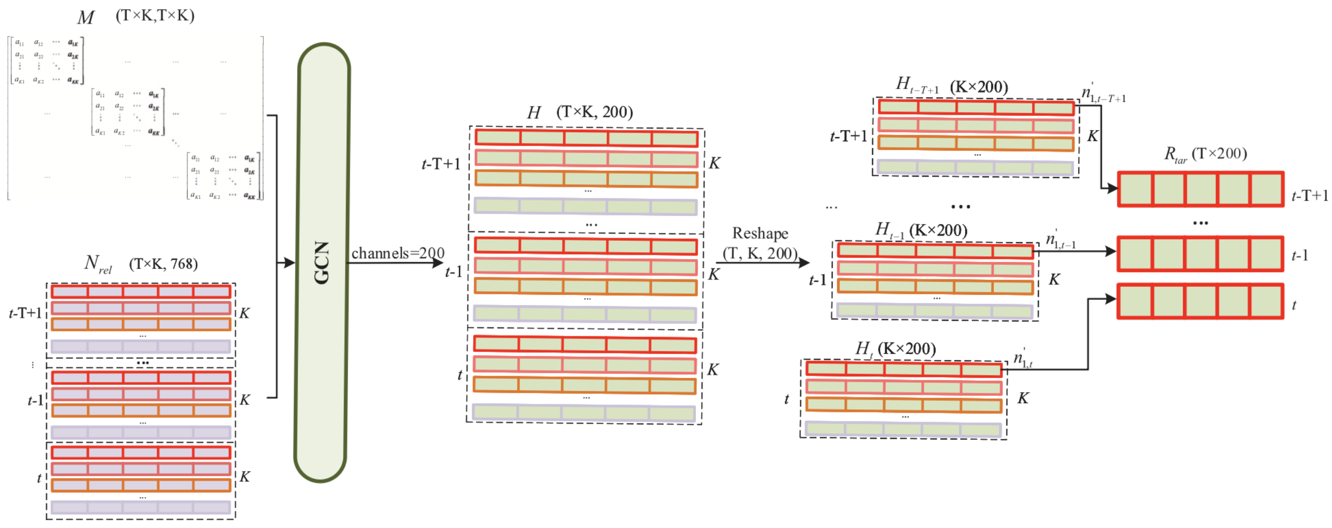

FIGURE 5. Fig. 5: GCN-based fusion of related-stock information

where $h_{t+1} \in \mathbb{R}^{1\times256}$ is the concatenation of the last forward and backward hidden states of $BiLSTM_2(\cdot)$ after dropout, and is used to predict the stock price movement on day $t+1$. Finally, h_{t+1} is fed into a dense layer and a softmax layer to predict the probability of the movement label $\hat{y}_{t+1}$:

$$\hat{y}_{t+1} = softmax(Dense(h_{t+1})). \tag{20}$$

Cross-entropy loss is used to update the neural network:

$$\mathcal{L} = CrossEntropy(\hat{y}_{t+1}, y_{t+1}), \tag{21}$$

where y_{t+1} is the ground-truth label, namely upward or downward.

IV. EXPERIMENTS

A. DATASET

All stocks used in the experiments are from the Chinese stock market. We select six representative target stocks from different industries, none of which overlap with the pretraining data. These stocks have relatively large market capitalization within their industries and a relatively large number of financial news items. Basic information about the target stocks is shown in Table 3. We collect a total of 54,969 news headlines from January 2, 2018 to June 18, 2021, among which 11,040 are related to the six target stocks and the remaining 43,929 are related to their corresponding related stocks. Statistics of the training set (80%), validation set (10%), and test set (10%) are shown in Table 4. Following previous studies [14,63], we adopt walk-forward testing [63] to train and validate the model so as to make fuller use of the available data.

TABLE 3. Basic information about the target stocks.

Stock Code	Stock Name	Industry
000063	ZTE	Communication Equipment Manufacturing
000651	Gree Electric Appliances	Electrical Machinery Manufacturing
601800	China Communications Construction	Civil Engineering Construction
000876	New Hope	Agro-food Processing
600104	SAIC Motor	Automobile Manufacturing
601933	Yonghui Superstores	Wholesale and Retail

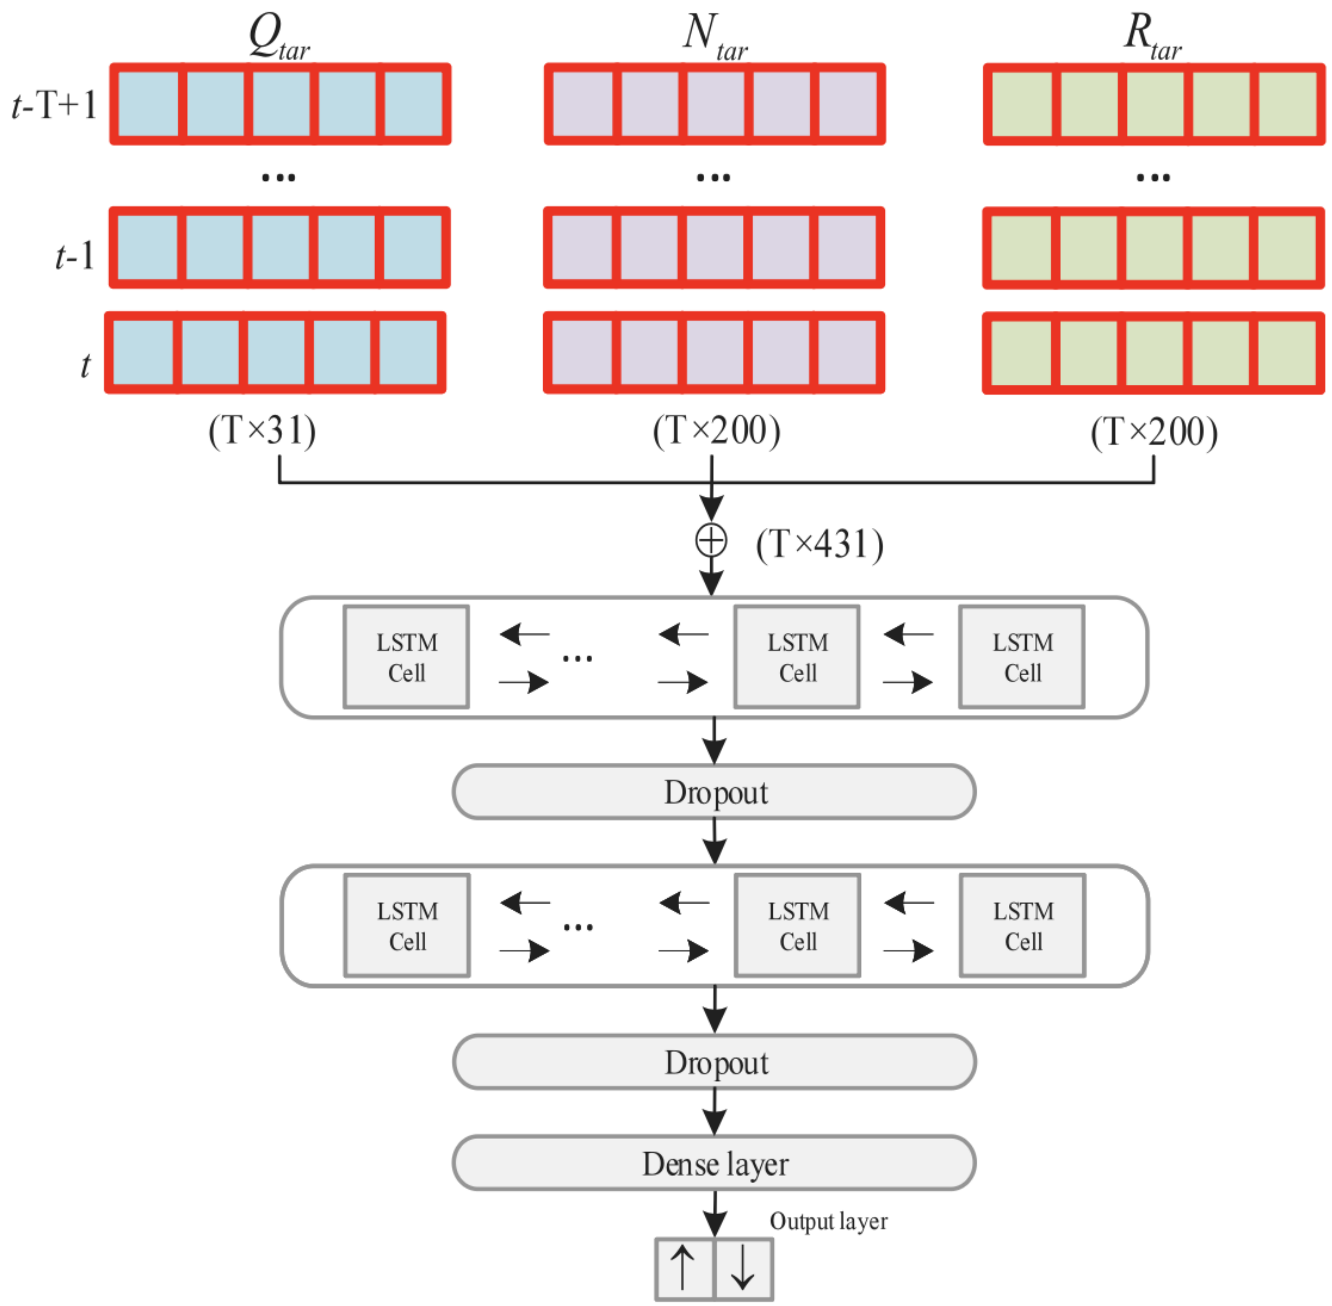

FIGURE 6. **Fig. 6: Stock price movement prediction based on BiLSTM**

B. BASELINES

To evaluate the performance of the proposed method, we compare it with several baselines, including eLSTM, LSTM-news, GCN-S, and RoBERTa. The baseline methods are summarized as follows:

- eLSTM [17] is a tensor-based event-driven LSTM model that can effectively use news signals for stock price movement prediction. The original authors built a lexicon to generate sentiment features; we follow their original settings and use the lexicon they constructed.
- LSTM-news [14] is an LSTM-based stock prediction model that combines technical indicators with target-stock news sentiment. In the original work, the news sentiment features were derived from an English financial sentiment lexicon. Because there is no corresponding Chinese lexicon, we use the pretrained MSC to generate news features, while following the original settings for the quantitative indicators and model structure.

TABLE 4. Dataset statistics.

Stock	Dataset	Start Date	End Date	# tar	# rel	# days	# up	# down
000063	Train	02/01/2018	21/10/2020	1992	3924	640	325	315
000063	Validation	22/10/2020	18/02/2021	540	2157	80	37	43
000063	Test	19/02/2021	18/06/2021	1011	2732	80	37	43
000063	All	02/01/2018	18/06/2021	3543	8813	800	399	401
000651	Train	02/01/2018	14/10/2020	1422	2054	668	329	339
000651	Validation	15/10/2020	09/02/2021	404	1547	83	46	37
000651	Test	10/02/2021	18/06/2021	784	2153	83	40	43
000651	All	02/01/2018	18/06/2021	2610	5754	834	415	419
601800	Train	02/01/2018	12/10/2020	209	3850	672	315	357
601800	Validation	13/10/2020	08/02/2021	92	1424	84	32	52
601800	Test	09/02/2021	18/06/2021	67	2162	84	40	44
601800	All	02/01/2018	18/06/2021	368	7436	840	387	453
000876	Train	02/01/2018	12/10/2020	623	2530	672	349	323
000876	Validation	13/10/2020	08/02/2021	429	1561	84	40	44
000876	Test	09/02/2021	18/06/2021	569	2259	84	31	53
000876	All	02/01/2018	18/06/2021	1621	6350	840	420	420
600104	Train	02/01/2018	12/10/2020	863	4474	672	331	341
600104	Validation	13/10/2020	08/02/2021	464	2315	84	40	44
600104	Test	09/02/2021	18/06/2021	695	3578	84	40	44
600104	All	02/01/2018	18/06/2021	2022	10367	840	411	429
601933	Train	02/01/2018	13/10/2020	399	1927	672	329	343
601933	Validation	14/10/2020	09/02/2021	203	1221	84	35	49
601933	Test	10/02/2021	18/06/2021	274	2061	83	34	49
601933	All	02/01/2018	18/06/2021	876	5209	839	398	441
Overall	All	02/01/2018	18/06/2021	11040	43929	4993	2430	2563

Note: `# tar` denotes the number of target-stock news items, `# rel` denotes the number of related-stock news items, `# days` denotes the number of trading days, `# up` denotes the number of upward days, and `# down` denotes the number of downward days.

- GCN-S [28] is a GCN-based stock price movement prediction model that predicts stock movement by incorporating a shareholding graph. The features used are OP, CP, HP, LP, and Vol (see Table 2). We adopt the original settings of GCN-S.
- RoBERTa [33] is a widely used pretrained text representation model. We use pretrained RoBERTa to extract news features. In this baseline, the target-stock news feature on trading day t is directly used to predict the stock price movement on trading day $t+1$.

Because our experiments use Chinese news data and we do not employ a version of FinBERT adapted to Chinese financial text, FinBERT [64] is not included in the comparison.

C. EXPERIMENTAL SETTINGS

The basic framework of the model is implemented in Keras and TensorFlow. The batch size is set to 32, and the maximum number of training epochs is set to 200. Early stopping is adopted to automatically terminate training and reduce the risk of overfitting. The model parameters are trained using the Adam optimizer with an initial learning rate of 0.001. The hidden-state dimension of BiLSTM is set to 128, and activity regularization L2(0.001) is applied to the kernel weight matrix. In addition, we further investigate several key hyperparameters on the validation set, including the time step T (ranging over $\{1, 2, ..., 14\}$), the number of related stocks K (ranging over $\{1, 2, ..., 15\}$), and the number of BiLSTM layers (ranging over $\{1, 2, 3, 4\}$), which are analyzed in detail in Section 5.4. The comparative results on the test set are all based on the parameter setting that performs well on the validation set, namely time step $T = 6$, number of related stocks $K = 10$, and two BiLSTM layers.

D. EVALUATION METRICS

Evaluating data science models in the financial domain is relatively complex [44]. Therefore, we assess the proposed model from both the classification and financial perspectives. For classification performance, following previous studies [17,49,50], we use accuracy (ACC) and Matthews correlation coefficient (MCC) as evaluation metrics. MCC can effectively mitigate the bias caused by data imbalance. The two metrics are defined as follows:

$$ACC = \frac{TP + TN}{TP + TN + FP + FN}, \tag{22}$$

$$MCC = \frac{TP \times TN - FP \times FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}}, \tag{23}$$

where TP is true positive, FP is false positive, TN is true negative, and FN is false negative. For financial evaluation, we simulate stock trading according to the predicted stock price movements, focusing on return and risk. We use Total Money (TMoney) [65] and the Sharpe ratio to evaluate simulated trading performance. The Sharpe ratio measures excess return per unit of volatility or total risk. The annual risk-free rate is set to 3.0%. These two metrics are defined as follows:

$$TMoney = \text{Available Capital} + \text{Closing Price} \times \text{Number of Stocks}, \tag{24}$$

$$\text{Sharpe Ratio} = \frac{\text{Rate of Return} - \text{Risk Free Rate}}{\text{Standard Deviation of Return}}. \tag{25}$$

V. RESULTS

A. STOCK PRICE MOVEMENT PREDICTION RESULTS

In this section, we compare the classification performance of MAC with that of the compared baselines. The results are shown in Table 5. Overall, the proposed method achieves the highest ACC and MCC on all six target stocks, improving on average by 2.38% and 4.62%, respectively, over eLSTM, the strongest baseline. These results indicate that MAC has a relatively stable advantage in classification performance.

Several factors may explain these results. First, compared with eLSTM, MAC uses a GCN module to model the structural information of the stock relation graph, which allows it to aggregate the influence of news from related stocks more effectively. In addition, market-driven sentiment classification pretraining helps improve the discriminative ability of news features. Second, compared with LSTM-news, which uses only target-stock quantitative indicators and news features, the proposed method improves ACC and MCC by 7.81% and 15.88%, respectively, indicating that news about related companies provides effective complementary information for target-stock movement prediction. Third, compared with GCN-S, which uses only a single shareholding graph and price features, MAC achieves better results, suggesting that the stock relation modeling strategy and news feature extraction method adopted in this paper are more effective. As for RoBERTa, although its results show that pretrained language models help represent news, there is still a noticeable gap between RoBERTa and MAC. This further demonstrates the necessity of jointly modeling temporal information and multi-source features in the financial domain.

In summary, incorporating temporal dependence in financial data, market-driven sentiment pretraining, and multi-source information fusion helps improve stock price movement prediction.

TABLE 5. Results of different prediction models.

Method	Metric	000063	000651	601800	000876	600104	601933	Average
RoBERTa	ACC	0.6017	0.5941	0.5688	0.5652	0.5909	0.5962	0.5862
RoBERTa	MCC	0.1949	0.1855	0.0899	0.0685	0.1750	0.1744	0.1480
GCN-S	ACC	0.5676	0.5613	0.5962	0.5962	0.5897	0.5897	0.5835
GCN-S	MCC	0.1471	0.1272	0.2026	0.1348	0.1759	0.1136	0.1502
LSTM-news	ACC	0.6096	0.6078	0.5714	0.5897	0.5909	0.5974	0.5945
LSTM-news	MCC	0.2174	0.2717	0.1144	0.1524	0.1840	0.1782	0.1864
eLSTM	ACC	0.6795	0.6383	0.6646	0.6218	0.6667	0.6220	0.6488
eLSTM	MCC	0.3548	0.3160	0.3147	0.2347	0.3326	0.2414	0.2990
MAC (Ours)	ACC	0.7000	0.6828	0.6859	0.6346	0.6923	0.6402	0.6726
MAC (Ours)	MCC	0.4017	0.3697	0.3593	0.2775	0.3821	0.2808	0.3452

B. ABLATION STUDY

To analyze the effectiveness of each feature component in MAC, we conduct ablation experiments with the following settings: (A) quantitative indicators of the target stock (Q_{tar}), (B) news features of the target stock (N_{tar}), and (C) news features of related stocks (R_{tar}). Here, A+B+C denotes the complete MAC model, while A, B, C, A+B, and A+C are used as ablation baselines. The news features in B and C are generated by the pretrained MSC module. To further verify the role of MSC and the pretraining strategy, we also introduce two additional news-feature generation methods for comparison: A+B_0+C_0 indicates that the news features are extracted by the original Chinese RoBERTa without domain pretraining, whereas A+B_1+C_1 indicates that the news features are extracted by Chinese RoBERTa pretrained on a semantics-based financial news sentiment analysis dataset. The ablation results are shown in Table 6.

According to the results, when only a single feature source is used, the news features (B and C) outperform the quantitative indicators (A), indicating that news information has stronger discriminative power than pure numerical information in this task. When two feature types are fused, the model performance further improves, and A+B performs better than A+C, suggesting that the target-stock news and quantitative indicators are more complementary. When all three feature types are included, A+B+C achieves the highest ACC and MCC, indicating that the full model can integrate multi-source information more effectively.

Comparing different pretraining strategies, among A+B_0+C_0, A+B_1+C_1, and A+B+C, the original Chinese RoBERTa without domain pretraining (A+B_0+C_0) yields the lowest average results. Chinese RoBERTa pretrained on the semantics-driven sentiment analysis dataset (A+B_1+C_1) improves over the original RoBERTa, increasing ACC and MCC by 0.77% and 1.00%, respectively. The proposed market-driven sentiment classification pretraining (A+B+C) achieves still better results, improving ACC and MCC by 1.94% and 3.71%, respectively, over A+B_1+C_1. This is because market-driven sentiment classification pretraining corresponds more directly to the actual stock market response to news than semantics-based sentiment polarity analysis does, and is therefore better suited to stock price movement prediction.

In summary, both the complete multi-source feature fusion structure and the market-driven pretraining mechanism are important contributors to the results achieved by MAC.

TABLE 6. Ablation results.

Features	Metric	000063	000651	601800	000876	600104	601933	Average
A (Q_{tar})	ACC	0.5779	0.5652	0.5803	0.5803	0.5741	0.5741	0.5753
A (Q_{tar})	MCC	0.1531	0.1274	0.1274	0.1274	0.1408	0.0864	0.1191
B (N_{tar})	ACC	0.5890	0.6144	0.5741	0.5897	0.5769	0.5962	0.5901
B (N_{tar})	MCC	0.1701	0.2682	0.0903	0.1065	0.1505	0.1744	0.1600
C (R_{tar})	ACC	0.5867	0.6138	0.5786	0.5833	0.5833	0.5976	0.5906
C (R_{tar})	MCC	0.1656	0.2294	0.1218	0.1161	0.1560	0.1250	0.1523
A+B	ACC	0.6313	0.6340	0.5926	0.6169	0.6234	0.6173	0.6193
A+B	MCC	0.2560	0.3090	0.1660	0.1831	0.2405	0.2219	0.2294
A+C	ACC	0.6233	0.6340	0.5926	0.6169	0.6115	0.6013	0.6133
A+C	MCC	0.2438	0.3090	0.1441	0.1831	0.2188	0.1593	0.2097
A+B_0+C_0	ACC	0.6800	0.6621	0.6539	0.6218	0.6795	0.6220	0.6532
A+B_0+C_0	MCC	0.3671	0.3281	0.2986	0.2710	0.3583	0.2252	0.3081
A+B_1+C_1	ACC	0.6953	0.6690	0.6654	0.6218	0.6859	0.6282	0.6609
A+B_1+C_1	MCC	0.3872	0.3432	0.3165	0.2710	0.3706	0.2202	0.3181
A+B+C	ACC	0.7000	0.6828	0.6859	0.6346	0.6923	0.6402	0.6726
A+B+C	MCC	0.4017	0.3697	0.3593	0.2775	0.3821	0.2808	0.3452

Note: A denotes quantitative indicators of the target stock (Q_{tar}); B denotes target-stock news features (N_{tar}); C denotes related-stock news features (R_{tar}); $A + B_0 + C_0$ denotes the concatenation of all features, where the news features come from the original Chinese RoBERTa without domain pretraining; $A + B_1 + C_1$ denotes the concatenation of all features, where the news features come from Chinese RoBERTa pretrained on a semantics-based financial news sentiment analysis dataset; and $A + B + C$ denotes the complete MAC model proposed in this paper.

C. FINANCIAL EVALUATION

For financial evaluation, we simulate stock trading (backtesting) based on the stock price movement predictions of different models. The backtesting data are formed by combining the validation set and test set of each target stock. We also adopt the buy-and-hold (B&H) strategy as an additional financial benchmark. The initial investment capital is set to RMB 10,000, and transaction costs are assumed to be zero throughout backtesting. If the predicted label for the next trading day ($t + 1$) is 1 (upward), the target stock is bought at the closing price on the current trading day (t), and all available capital is invested. If the predicted label is 0 (downward), the stock is sold at the closing price on day t. If the predicted label for day $t + 1$ is the same as the label predicted on the previous day for the current day t, no trading action is taken. Under the buy-and-hold strategy, all capital is invested at the initial time and the stock is held throughout the period without subsequent trading [68,69]. Therefore, changes in capital under the buy-and-hold strategy also reflect the stock price trend in real time.

The backtesting results are shown in Fig. 7. Overall, compared with the baselines, the MAC model achieves higher returns on all target stocks. Except for stock 601933, the proposed method yields positive returns on the other five target stocks. For stock 601933, although the return remains negative, the loss of MAC is smaller than those of the baselines, and MAC yields RMB 745.96 more than eLSTM. This indicates that MAC not only attains higher ACC in the classification task (see Section 5.1), but also captures trading signals more effectively and converts them into actual returns.

Furthermore, Table 7 reports the Sharpe ratios of different methods. A higher Sharpe ratio indicates higher return per unit of risk. As shown in Table 7, MAC achieves better Sharpe ratios on all six target stocks, which indicates that the proposed method performs better not only in terms of return, but also in terms of the trade-off between return and risk. The negative Sharpe ratio for stock 601933 means that its return does not exceed the risk-free rate of 3.0%. This phenomenon is largely related to the substantial decline in the market performance of that stock, as shown by the capital curve of the buy-and-hold strategy in Fig. 7. Even so, MAC still shows a relative advantage over the other baselines.

In summary, the financial evaluation further indicates that MAC can not only improve classification accuracy, but also achieve higher returns and better risk-return balance in practical trading scenarios.

TABLE 7. Sharpe ratios of different methods.

Model	000063	000651	601800	000876	600104	601933
B&H	0.0065	-1.1768	-0.8146	-2.0120	-0.0249	-2.5294
GCN-S	1.0378	-0.4178	0.8530	-0.0029	1.5067	-1.7214
RoBERTa	1.2566	-0.0382	0.8782	0.1258	1.2418	-1.7184
LSTM-news	1.3609	0.1043	0.9847	0.1969	1.2621	-1.9863
eLSTM	2.3786	0.2975	1.6989	0.5676	1.8346	-0.9372
MAC (Ours)	3.0053	1.0384	2.1627	0.9668	2.8595	-0.1238

D. HYPERPARAMETER ANALYSIS

Because financial data are inherently sequential, hyperparameter settings have an important impact on model performance. To capture temporal financial information more effectively, we investigate the effects of the window size $T \in \{1, 2, ..., 14\}$, the number of related stocks $K \in \{1, 2, ..., 15\}$, and the number of BiLSTM layers (1, 2, 3, and 4) on MAC. All related experiments are conducted on the validation sets of the six target stocks.

The results are shown in Figs. 8 and 9. As can be seen from Fig. 8, most stocks achieve relatively high ACC when $T = 6$ and $K = 10$, indicating that a moderate time window and a moderate number of related stocks are more beneficial for model learning. Further increasing T and K does not improve model performance, which suggests that an overly long historical window and too many related stocks may introduce additional noise and consequently weaken prediction performance. As shown in Fig. 9, when BiLSTM is configured with two layers, the model achieves the best validation performance, with the highest average ACC and MCC across the six stocks. Too few layers may be insufficient to capture temporal patterns, whereas too many layers may increase model complexity and hurt generalization. Accordingly, we finally adopt $T = 6$, $K = 10$, and two BiLSTM layers as the hyperparameter configuration used in this study.

E. FINE-GRAINED ANALYSIS

As described in Section 3.2.1, listed-company news can be divided into three categories: company information, business information, and financial information. Meanwhile, news can also be divided by sentiment polarity into positive news associated with stock increases and negative news associated with stock decreases. To analyze the effects of different types of news on stock price movement prediction, we conduct fine-grained experiments based on these categories and jointly consider each type of news feature with quantitative indicators. The results are shown in Table 8.

According to the results, compared with the models that incorporate only company-information news (A+$B_{\mathcal{C}}$+$C_{\mathcal{C}}$) or business-information news (A+$B_{\mathcal{B}}$+$C_{\mathcal{B}}$), the model that incorporates financial-information news (A+$B_{\mathcal{F}}$+$C_{\mathcal{F}}$) performs better in both ACC and MCC, suggesting that financial news is more directly related to stock price movements. On the other hand, the model based on negative news (A+$B_{\mathcal{N}}$+$C_{\mathcal{N}}$) outperforms the model based on positive news (A+$B_{\mathcal{P}}$+$C_{\mathcal{P}}$) by 1.12% and 1.81% in average ACC and MCC, respectively, indicating that negative news usually provides stronger signals for stock price movement prediction. This phenomenon may be related to the market being more sensitive to negative information.

Furthermore, when all types of news are jointly considered in MAC (A+B+C), the model achieves the highest ACC and MCC. This result indicates that different categories of news have a certain degree of complementarity, and that their joint use can provide a more comprehensive representation of information affecting stock price movement.

In summary, different types of news contribute differently to stock price movement prediction. Among them, financial news and negative news appear to be more predictive, while jointly exploiting multiple types of news yields the best overall results.

VI. CONCLUSION

To capture the influence of news from related stocks, this paper proposes a model that combines GCN and BiLSTM to predict stock price movements in the Chinese stock market. The model integrates multi-source information, including quantitative indicators, news about the target company, and news about companies related to the target company. To obtain more effective news embeddings, we pretrain a market-driven sentiment classifier to reflect the responses of the market and investors to news.

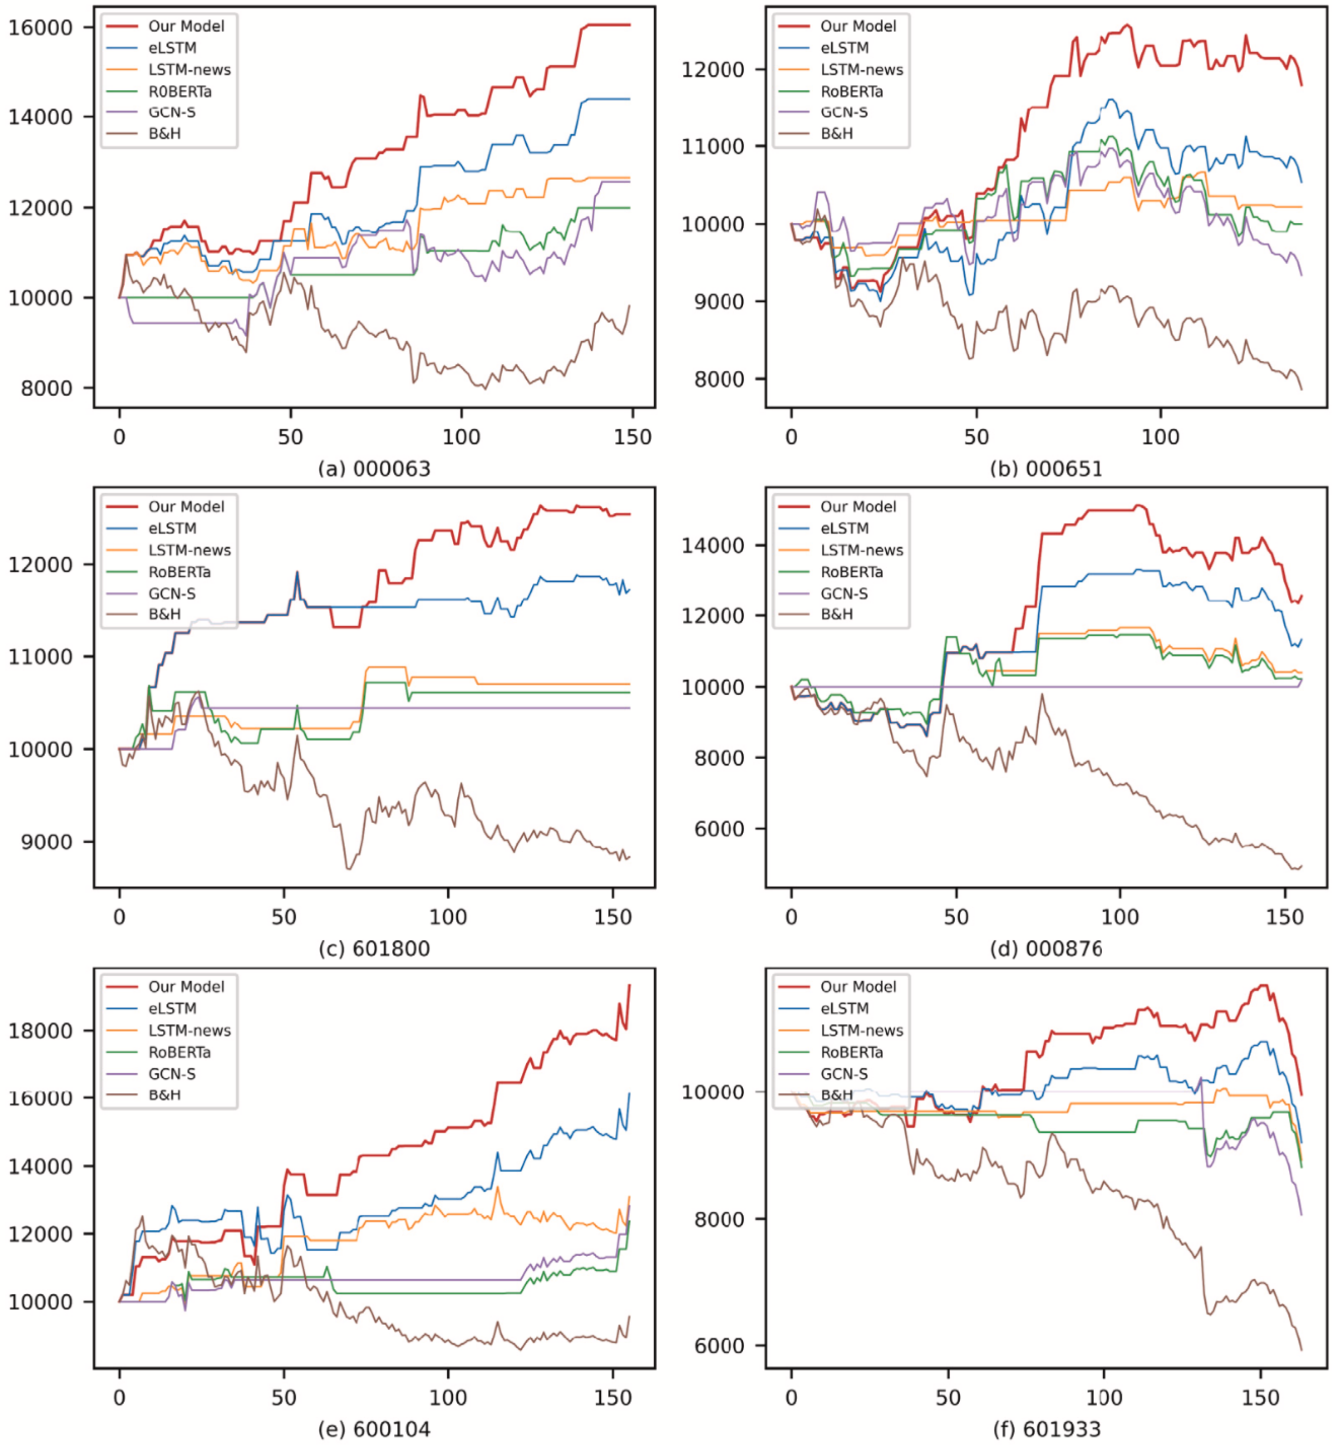

FIGURE 7. **Fig. 7: TMoney of different models**

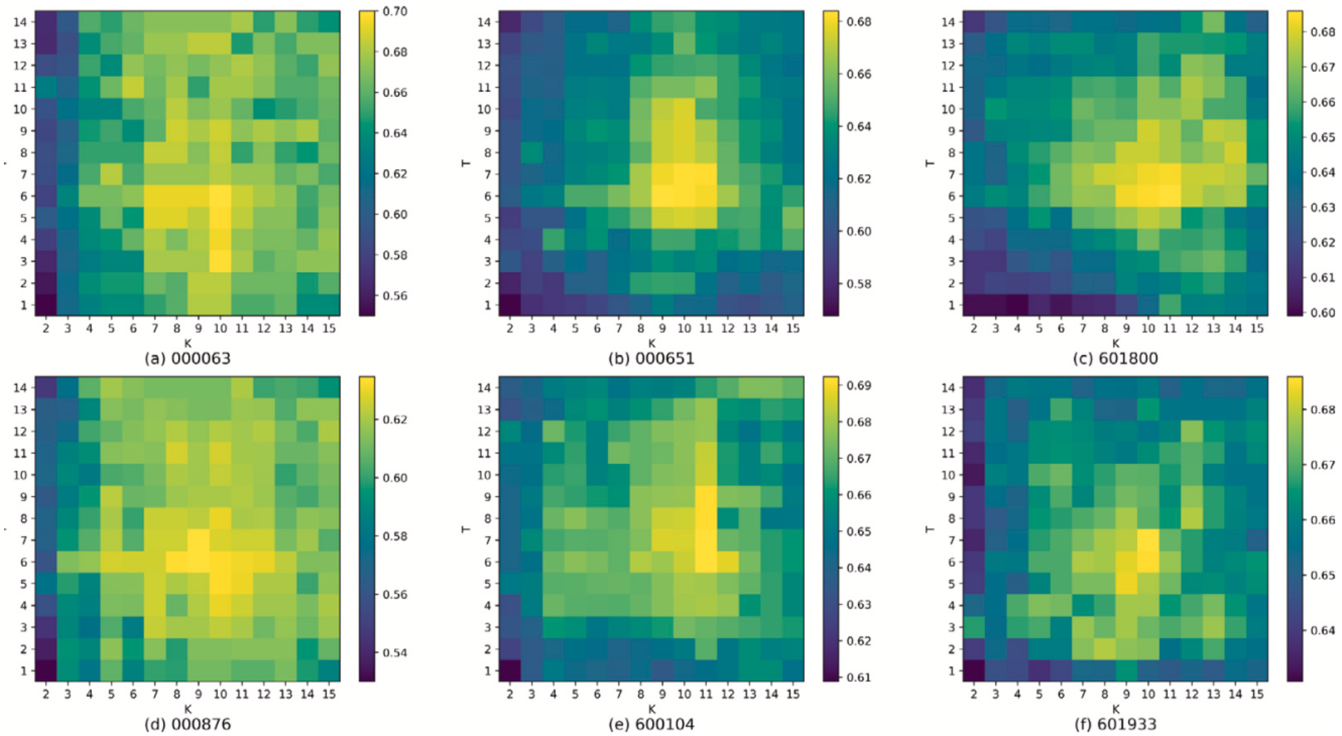

FIGURE 8. **Fig. 8: Prediction ACC under different window sizes and numbers of related stocks**

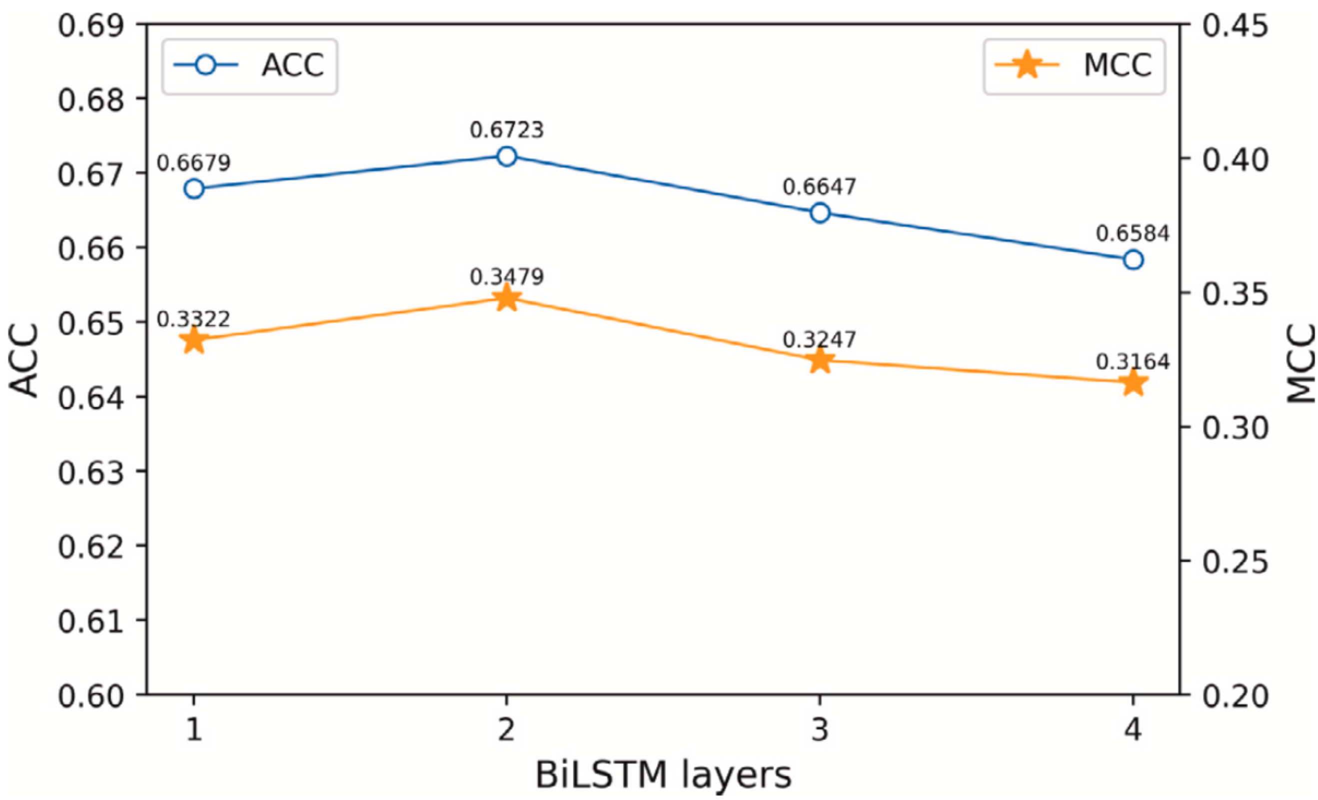

FIGURE 9. **Fig. 9: Prediction ACC and MCC under different numbers of BiLSTM layers**

TABLE 8. **Fine-grained analysis results.**

Features	Metric	000063	000651	601800	000876	600104	601933	Average
A+$B_{\mathcal{C}}$+$C_{\mathcal{C}}$	ACC	0.6223	0.6306	0.6064	0.6017	0.6239	0.6098	0.6158
A+$B_{\mathcal{C}}$+$C_{\mathcal{C}}$	MCC	0.2410	0.2617	0.1684	0.1949	0.2468	0.1963	0.2182
A+$B_{\mathcal{B}}$+$C_{\mathcal{B}}$	ACC	0.6310	0.6343	0.6104	0.6096	0.6330	0.6037	0.6203
A+$B_{\mathcal{B}}$+$C_{\mathcal{B}}$	MCC	0.2577	0.2675	0.1786	0.2152	0.2590	0.2155	0.2323
A+$B_{\mathcal{F}}$+$C_{\mathcal{F}}$	ACC	0.6388	0.6370	0.6154	0.6104	0.6463	0.6159	0.6273
A+$B_{\mathcal{F}}$+$C_{\mathcal{F}}$	MCC	0.2717	0.2706	0.1911	0.2168	0.2916	0.1962	0.2397
A+$B_{\mathcal{P}}$+$C_{\mathcal{P}}$	ACC	0.6407	0.6403	0.6096	0.6090	0.6296	0.6154	0.6241
A+$B_{\mathcal{P}}$+$C_{\mathcal{P}}$	MCC	0.2788	0.2944	0.1782	0.1904	0.2634	0.2508	0.2427
A+$B_{\mathcal{N}}$+$C_{\mathcal{N}}$	ACC	0.6543	0.6414	0.6194	0.6098	0.6667	0.6201	0.6353
A+$B_{\mathcal{N}}$+$C_{\mathcal{N}}$	MCC	0.3080	0.2850	0.2144	0.1868	0.3309	0.2399	0.2608
A+B+C	ACC	0.7000	0.6828	0.6859	0.6346	0.6923	0.6402	0.6726
A+B+C	MCC	0.4017	0.3697	0.3593	0.2775	0.3821	0.2808	0.3452

Note: $A+B+C$ denotes the complete MAC model proposed in this paper, which uses quantitative indicators of the target stock (A), all types of news about the target stock (B), and all types of news about related stocks (C). $A+B_{\mathcal{C}}+C_{\mathcal{C}}$, $A+B_{\mathcal{B}}+C_{\mathcal{B}}$, and $A+B_{\mathcal{F}}+C_{\mathcal{F}}$ use company-information, business-information, and financial-information news, respectively. $A+B_{\mathcal{P}}+C_{\mathcal{P}}$ and $A+B_{\mathcal{N}}+C_{\mathcal{N}}$ use positive news and negative news, respectively.

We validate the proposed method on six target stocks from different industries. The experimental results show that the proposed method improves ACC and MCC by 2.38% and 4.62%, respectively, over the strongest baseline in our comparative experiments. In financial evaluation, the model also performs better, yielding higher returns and higher Sharpe ratios. These results indicate that aggregating news features from related stocks helps improve prediction performance. In addition, because the market-driven sentiment classifier corresponds more directly to actual market responses to news, it outperforms semantics-based sentiment polarity analysis in news feature embedding. Overall, the proposed method can improve stock price movement prediction and provide more effective decision support for investors seeking to optimize investment strategies and reduce the risks associated with margin trading and securities lending.

The model still has several limitations that should be addressed in future work. First, it cannot determine the authenticity of news, and some news about the target stock may be false information released by competitors for strategic purposes. Future work may introduce fake-news detectors [68] and further examine the differential effects of real and fake news on the stock market. Second, financial news often contains figurative language, such as metaphor, which poses challenges to computational language understanding. To address this issue, metaphor processing techniques [69,70] may be incorporated to convert metaphorical expressions into more machine-processable literal expressions, thereby further improving stock price movement prediction.

REFERENCES

REFERENCES

[1] W. Jiang, Applications of deep learning in stock market prediction: Recent progress, Expert Syst. Appl. 184 (2021) 115537, https://doi.org/10.1016/j.eswa.2021.115537.

[2] O. Bustos, A. Pomares-Quimbaya, Stock market movement forecast: A systematic review, Expert Syst. Appl. 156 (2020) 113464, https://doi.org/10.1016/j.eswa.2020.113464.

[3] S. Merello, A. Picasso, L. Oneto, E. Cambria, Ensemble application of transfer learning and sample weighting for stock market prediction, in: IJCNN, 2019.

[4] Q. Li, Y. Chen, L.L. Jiang, P. Li, H. Chen, A tensor-based information framework for predicting the stock market, ACM Trans. Inf. Syst. 34 (2) (2016) https://doi.org/10.1145/2838731.

[5] F. Xing, L. Malandri, Y. Zhang, E. Cambria, Financial sentiment analysis: An investigation into common mistakes and silver bullets, in: COLING, 2020, pp. 978-987.

[6] I.K. Nti, A.F. Adekoya, B.A. Weyori, A systematic review of fundamental and technical analysis of stock market predictions, Artif. Intell. Rev. 53 (4) (2020) 3007-3057, https://doi.org/10.1007/s10462-019-09754-z.

[7] Q. Li, Y. Chen, J. Wang, Y. Chen, H. Chen, Web media and stock markets: A survey and future directions from a big data perspective, IEEE Trans. Knowl. Data Eng. 30 (2) (2018) 381-399, https://doi.org/10.1109/TKDE.2017.2763144.

[8] P.C. Tetlock, M. Saar-Tsechansky, S. Macskassy, More than words: Quantifying language to measure firms' fundamentals, J. Financ. 63 (3) (2008) 1437-1467, https://doi.org/10.1111/j.1540-6261.2008.01362.x.

[9] Y. Jeon, T.H. McCurdy, X. Zhao, News as sources of jumps in stock returns: Evidence from 21 million news articles for 9000 companies, J. Financ. Econ. (2021) https://doi.org/10.1016/j.jfineco.2021.08.002.

[10] Y. Shi, W. Li, L. Zhu, K. Guo, E. Cambria, Stock trading rule discovery with double deep Q-network, Appl. Soft Comput. 107 (107320) (2021).

[11] Q. Li, T. Wang, P. Li, L. Liu, Q. Gong, Y. Chen, The effect of news and public mood on stock movements, Inform. Sci. 278 (2014) 826-840, https://doi.org/10.1016/j.ins.2014.03.096.

[12] L. Malandri, F. Xing, C. Orsenigo, C. Vercellis, E. Cambria, Public mood-driven asset allocation: the importance of financial sentiment in portfolio management, Cogn. Comput. 10 (6) (2018) 1167-1176.

[13] A. Thakkar, K. Chaudhari, Fusion in stock market prediction: A decade survey on the necessity, recent developments, and potential future directions, Inf. Fusion 65 (2021) 95-107, https://doi.org/10.1016/j.inffus.2020.08.019.

[14] X. Li, P. Wu, W. Wang, Incorporating stock prices and news sentiments for stock market prediction: A case of Hong Kong, Inf. Process. Manage. 57 (5) (2020) 102212, https://doi.org/10.1016/j.ipm.2020.102212.
[15] W. Chen, C.K. Yeo, C.T. Lau, B.S. Lee, Leveraging social media news to predict stock index movement using RNN-boost, Data Knowl. Eng. 118 (2018) 14-24, https://doi.org/10.1016/j.datak.2018.08.003.
[16] M.Y. Chen, C.H. Liao, R.P. Hsieh, Modeling public mood and emotion: Stock market trend prediction with anticipatory computing approach, Comput. Hum. Behav. 101 (2019) 402-408, https://doi.org/10.1016/j.chb.2019.03.021.
[17] Q. Li, J. Tan, J. Wang, H. Chen, A multimodal event-driven LSTM model for stock prediction using online news, IEEE Trans. Knowl. Data Eng. 33 (10) (2021) 3323-3337, https://doi.org/10.1109/TKDE.2020.2968894.
[18] S. Merello, A. Picasso, Y. Ma, L. Oneto, E. Cambria, Investigating timing and impact of news on the stock market, in: Proceedings of ICDM Workshops, 2018, pp. 1348-1354.
[19] F. Xing, F. Pallucchini, E. Cambria, Cognitive-inspired domain adaptation of sentiment lexicons, Information Processing and Management 56 (3) (2019) 554-564.
[20] W. Li, L. Zhu, Y. Shi, K. Guo, E. Cambria, User reviews: sentiment analysis using lexicon integrated two-channel CNN-LSTM family models, Applied Soft Computing 94 (106435) (2020).
[21] A. Khadjeh Nassirtoussi, S. Aghabozorgi, T. Ying Wah, D.C.L. Ngo, Text mining for market prediction: A systematic review, Expert Syst. Appl. 41 (16) (2014) 7653-7670, https://doi.org/10.1016/j.eswa.2014.06.009.
[22] E. Cambria, Q. Liu, S. Decherchi, F. Xing, K. Kwok, SenticNet 7: A commonsense-based neurosymbolic AI framework for explainable sentiment analysis, in: Proceedings of the 13th Language Resources and Evaluation Conference, European Language Resources Association, 2022, pp. 3829-3839.
[23] A.E. de Oliveira Carosia, G.P. Coelho, A.E.A. da Silva, Investment strategies applied to the Brazilian stock market: A methodology based on sentiment analysis with deep learning, Expert Syst. Appl. 184 (2021) 115470, https://doi.org/10.1016/j.eswa.2021.115470.
[24] M.G. Sousa, K. Sakiyama, L.d.S. Rodrigues, P.H. Moraes, E.R. Fernandes, E.T. Matsubara, BERT for stock market sentiment analysis, in: 2019 IEEE 31st International Conference on Tools with Artificial Intelligence, ICTAI, 2019, pp. 1597-1601, https://doi.org/10.1109/ICTAI.2019.00231.
[25] M. Van de Kauter, D. Breesch, V. Hoste, Fine-grained analysis of explicit and implicit sentiment in financial news articles, Expert Syst. Appl. 42 (11) (2015) 4999-5010, https://doi.org/10.1016/j.eswa.2015.02.007.
[26] N. Pröllochs, S. Feuerriegel, D. Neumann, Statistical inferences for polarity identification in natural language, PLoS One 13 (12) (2018) 1-21, https://doi.org/10.1371/journal.pone.0209323.
[27] S. Pandit, C.E. Wasley, T. Zach, Information externalities along the supply chain: The economic determinants of suppliers' stock price reaction to their customers' earnings announcements, Contemp. Account. Res. 28 (4) (2011) 1304-1343, https://doi.org/10.1111/j.1911-3846.2011.01092.x.
[28] Y. Chen, Z. Wei, X. Huang, Incorporating corporation relationship via graph convolutional neural networks for stock price prediction, in: Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM'18, Association for Computing Machinery, 2018, pp. 1655-1658, https://doi.org/10.1145/3269206.3269269.
[29] S. He, H. Rui, A.B. Whinston, Social media strategies in product-harm crises, Inf. Syst. Res. 29 (2) (2018) 362-380, https://doi.org/10.1287/isre.2017.0707.
[30] T. Yu, M. Sengul, R.H. Lester, Misery Loves company: The spread of negative impacts resulting from an organizational crisis, Acad. Manage. Rev. 33 (2) (2008) 452-472, https://doi.org/10.5465/amr.2008.31193499.
[31] S.B. Smart, C.J. Zutter, Fundamentals of Investing, fourteenth ed, Pearson, London, United Kingdom, 2019.
[32] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, V. Stoyanov, RoBERTa: A robustly optimized BERT pretraining approach, 2019, https://doi.org/10.48550/ARXIV.1907.11692.
[33] Z. Zhao, H. Chen, J. Zhang, X. Zhao, T. Liu, W. Lu, X. Chen, H. Deng, Q. Ju, X. Du, UER: An open-source toolkit for pre-training models, 2019, https://doi.org/10.48550/ARXIV.1909.05658.
[34] T.N. Kipf, M. Welling, Semi-supervised classification with graph convolutional networks, in: 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, 2017, pp. 24-26, https://doi.org/10.48550/ARXIV.1609.02907.
[35] A. Graves, J. Schmidhuber, Framewise phoneme classification with bidirectional LSTM and other neural network architectures, Neural Netw. 18 (5-6) (2005) 602-610, https://doi.org/10.1016/j.neunet.2005.06.042.
[36] S.A. Sharpe, Financial market imperfections, firm leverage, and the cyclicality of employment, Am. Econ. Rev. 84 (4) (1994) 1060-1074.
[37] E.F. Fama, The behavior of stock-market prices, J. Bus. 38 (1) (1965) 34-105.
[38] Y. Shi, K.Y. Ho, News sentiment and states of stock return volatility: Evidence from long memory and discrete choice models, Financ. Res. Lett. 38 (2021) 101446, https://doi.org/10.1016/j.frl.2020.101446.
[39] N. Strauss, R. Vliegenthart, P. Verhoeven, Lagging behind? Emotions in newspaper articles and stock market prices in the Netherlands, Publ. Relat. Rev. 42 (4) (2016) 548-555, https://doi.org/10.1016/j.pubrev.2016.03.010.
[40] J. Lei, N. Dawar, J. Lemmink, Negative spillover in brand portfolios: Exploring the antecedents of asymmetric effects, J. Mark. 72 (3) (2008) 111-123, https://doi.org/10.1509/JMKG.72.3.111.
[41] B.W. Jacobs, V.R. Singhal, Shareholder value effects of the Volkswagen emissions scandal on the automotive ecosystem, Prod. Oper. Manag. 29 (10) (2020) 2230-2251, https://doi.org/10.1111/poms.13228.
[42] D.Y. Joe, F.D. Oh, Spillover effects within business groups: The case of Korean chaebols, Manage. Sci. 64 (3) (2018) 1396-1412, https://doi.org/10.1287/mnsc.2016.2596.
[43] P. Zou, G. Li, How emerging market investors' value competitors' customer equity: Brand crisis spillover in China, J. Bus. Res. 69 (9) (2016) 3765-3771, https://doi.org/10.1016/j.jbusres.2015.12.068.
[44] F.Z. Xing, E. Cambria, R.E. Welsch, Natural language based financial forecasting: a survey, Artif. Intell. Rev. 50 (1) (2018) 49-73, https://doi.org/10.1007/s10462-017-9588-9.
[45] A. Picasso, S. Merello, Y. Ma, L. Oneto, E. Cambria, Technical analysis and sentiment embeddings for market trend prediction, Expert Syst. Appl. 135 (2019) 60-70, https://doi.org/10.1016/j.eswa.2019.06.014.
[46] S. Hochreiter, J. Schmidhuber, Long short-term memory, Neural Comput. 9 (8) (1997) 1735-1780.
[47] S. Minaee, N. Kalchbrenner, E. Cambria, N. Nikzad, M. Chenaghlu, J. Gao, Deep learning based text classification: A comprehensive review, ACM Comput. Surv. 54 (3) (2021) 62, https://doi.org/10.1145/3439726.
[48] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of deep bidirectional transformers for language understanding, 2018, https://doi.org/10.48550/ARXIV.1810.04805.
[49] X. Zhang, Y. Zhang, S. Wang, Y. Yao, B. Fang, P.S. Yu, Improving stock market prediction via heterogeneous inf. fusion, Knowl.-Based Syst. 143 (2018) 236-247, https://doi.org/10.1016/j.knosys.2017.12.025.
[50] H. Huang, X. Liu, Y. Zhang, C. Feng, News-driven stock prediction via noisy equity state representation, Neurocomputing 470 (2022) 66-75, https://doi.org/10.1016/j.neucom.2021.10.092.
[51] A.F. Adoma, N.-M. Henry, W. Chen, Comparative analyses of Bert, Roberta, Distilbert, and Xlnet for text-based emotion recognition, in: 2020 17th International Computer Conference on Wavelet Active Media Technology and Information Processing, ICCWAMTIP, 2020, pp. 117-121, https://doi.org/10.1109/ICCWAMTIP51612.2020.9317379.
[52] Q. Liu, X. Geng, H. Huang, T. Qin, J. Lu, D. Jiang, MGRC: An end-to-end multigranularity reading comprehension model for question answering, IEEE Trans. Neural Netw. Learn. Syst. (2021) 1-12, https://doi.org/10.1109/TNNLS.2021.3107029.

[53] R. Mao, X. Li, Bridging towers of multitask learning with a gating mechanism for aspect-based sentiment analysis and sequential metaphor identification, in: Proceedings of the 35th AAAI Conference on Artificial Intelligence, 2021, pp. 13534-13542.

[54] M. Ge, R. Mao, E. Cambria, Explainable metaphor identification inspired by conceptual metaphor theory, in: Proceedings of the 36th AAAI Conference on Artificial Intelligence, 2022, pp. 10681-10689.

[55] K. He, L. Yao, J. Zhang, Y. Li, C. Li, Construction of genealogical knowledge graphs from obituaries: Multitask neural network extraction system, J. Med. Internet Res. 23 (8) (2021) e25670, https://doi.org/10.2196/25670.

[56] R. Mao, Q. Liu, K. He, W. Li, E. Cambria, The biases of pre-trained language models: An empirical study on prompt-based sentiment analysis and emotion detection, IEEE Trans. Affect. Comput. (2022) 1-11, https://doi.org/10.1109/TAFFC.2022.3204972.

[57] A. Grover, J. Leskovec, Node2vec: Scalable feature learning for networks, in: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Association for Computing Machinery, New York, USA, 2016, pp. 855-864, https://doi.org/10.1145/2939672.2939754.

[58] Q. Lin, Y. Zhu, H. Lu, K. Shi, Z. Niu, Improving university faculty evaluations via multi-view knowledge graph, Futur. Gener. Comput. Syst. 117 (2021) 181-192, https://doi.org/10.1016/j.future.2020.11.021.

[59] Q. Lin, J. Liu, L. Zhang, Y. Pan, X. Hu, F. Xu, H. Zeng, Contrastive graph representations for logical formulas embedding, IEEE Trans. Knowl. Data Eng. (2021) https://doi.org/10.1109/TKDE.2021.3139333, 1-1.

[60] R. Mao, C. Lin, F. Guerin, End-to-end sequential metaphor identification inspired by linguistic theories, in: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 2019, pp. 3888-3898, https://doi.org/10.18653/v1/P19-1378.

[61] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, Y. Bengio, Learning phrase representations using RNN encoder-decoder for statistical machine translation, in: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP, 2014, pp. 1724-1734.

[62] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, Adv. Neural Inf. Process. Syst. 30 (2017).

[63] A.P. Ratto, S. Merello, L. Oneto, Y. Ma, L. Malandri, E. Cambria, Ensemble of technical analysis and machine learning for market trend prediction, in: 2018 IEEE Symposium Series on Computational Intelligence, SSCI, 2018, pp. 2090-2096, https://doi.org/10.1109/SSCI.2018.8628795.

[64] Y. Yang, M.C.S. Uy, A. Huang, FinBERT: A pretrained language model for financial communications, 2020, arXiv preprint arXiv:2006.08097.

[65] W. Chen, M. Jiang, W.-G. Zhang, Z. Chen, A novel graph convolutional feature based convolutional neural network for stock trend prediction, Inform. Sci. 556 (2021) 67-94, https://doi.org/10.1016/j.ins.2020.12.068.

[66] J. Li, M. Sun, Scalable term selection for text categorization, in: Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL), 2007, pp. 774-782.

[67] J. Li, M. Sun, X. Zhang, A comparison and semi-quantitative analysis of words and character-bigrams as features in Chinese text categorization, in: Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics, 2006, pp. 545-552.

[68] S. Han, F. Ciravegna, Rumour detection on social media for crisis management, in: ISCRAM, 2019.

[69] R. Mao, C. Lin, F. Guerin, Word embedding and WordNet based metaphor identification and interpretation, in: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2018, pp. 1222-1231.

[70] R. Mao, X. Li, M. Ge, E. Cambria, MetaPro: A computational metaphor processing model for text pre-processing, Inf. Fusion 86-87 (2022) 30-43, https://doi.org/10.1016/j.inffus.2022.06.002.