

LLMFactor: Extracting Profitable Factors through Prompts for Explainable Stock Movement Prediction

Meiyun Wang*
The University of Tokyo
omiun20@g.ecc.u-tokyo.ac.jp

Kiyoshi Izumi
The University of Tokyo
izumi@sys.t.u-tokyo.ac.jp

Hiroki Sakaji
Hokkaido University
sakaji@ist.hokudai.ac.jp

Abstract

Recently, Large Language Models (LLMs) have attracted significant attention for their exceptional performance across a broad range of tasks, particularly in text analysis. However, the finance sector presents a distinct challenge due to its dependence on time-series data for complex forecasting tasks. In this study, we introduce a novel framework called LLMFactor, which employs Sequential Knowledge-Guided Prompting (SKGP) to identify factors that influence stock movements using LLMs. Unlike previous methods that relied on keyphrases or sentiment analysis, this approach focuses on extracting factors more directly related to stock market dynamics, providing clear explanations for complex temporal changes. Our framework directs the LLMs to create background knowledge through a fill-in-the-blank strategy and then discerns potential factors affecting stock prices from related news. Guided by background knowledge and identified factors, we leverage historical stock prices in textual format to predict stock movement. An extensive evaluation of the LLMFactor framework across four benchmark datasets from both the U.S. and Chinese stock markets demonstrates its superiority over existing state-of-the-art methods and its effectiveness in financial time-series forecasting.

1 Introduction

Artificial Intelligence (AI) has become a staple in the financial sector, addressing various challenges such as predicting stock movements (Patel et al., 2015; Henrique et al., 2019), providing robo-advisory services (Xue et al., 2018; Bertrand et al., 2023), and managing risks (Ahbali et al., 2022; Wang et al., 2021). Among these tasks, forecasting stock trends is particularly significant, as it leverages historical data to shape trading strategies and pinpoint opportunities for buying or selling stocks.

*Corresponding Author

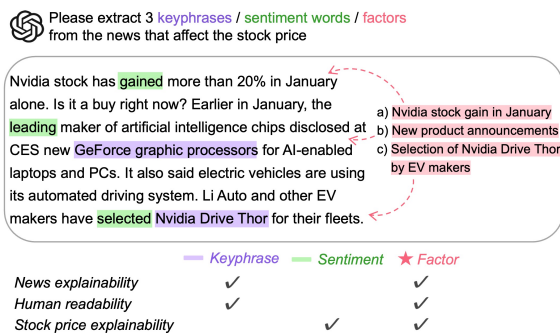


Figure 1: Comparing keyphrases, sentiment words, and factors: LLM-extracted factors offer greater human readability and explainability.

The Efficient Market Hypothesis (EMH), proposed by Eugene Fama (Fama, 1970), suggests that stock prices reflect all available information, making it difficult to predict future price movements. However, subsequent research has identified limitations to market efficiency, highlighting how phenomena such as information asymmetry (Timmermann and Granger, 2004) and irrational behaviours (Lo, 2004) can induce departures from perfect efficiency. These observations have paved the way for researchers to seek excess market returns by identifying and leveraging market inefficiencies.

Building on this foundation, there is a growing interest in exploring various data types to enhance prediction capabilities. Several studies underscore the importance of news related to stocks in uncovering fundamental market insights (Vargas et al., 2018; Li et al., 2021), while others highlight the significance of understanding the interconnections between companies and industries (Feng et al., 2019; Hsu et al., 2021). Recent research has empirically demonstrated the impact of public sentiment on market trends, with efforts to extract sentiment and keyphrases from news and social media data (Nguyen et al., 2015; Hao et al., 2021).

However, as illustrated in Figure 1, these meth-

ods encounter various limitations. We propose a novel task to enhance stock movement prediction through the use of “**factors.**” Figure 1 delineates the distinctions among keyphrases, sentiment words, and factors generated by LLMs. For the given stock-related news, keyphrases may outline the content but do not directly correlate with the stock price. Conversely, sentiments are associated with stock price but lack clarity for human interpretation. Consequently, factors offer superiority in three aspects: **news explainability, human readability, and stock price explainability.**

To incorporate factors into financial forecasting, we introduce a new framework **LLMFactor**, which derives factors from LLMs through **Sequential Knowledge-Guided Prompting (SKGP)** and then explains stock price trends. LLMFactor represents a holistic approach to integrating LLMs into financial applications. Initially, our proposed SKGP strategy prompts LLMs to generate factors related to stocks. SKGP starts with a fill-in-the-blank technique to elicit background knowledge (e.g. company relations) related to the stock. Subsequently, we prompt LLMs to identify reliable factors from news articles. Finally, we compile the text-formatted time-series data with these factors and knowledge to predict stock price trends. Experiments across four datasets indicate that LLMFactor yields superior predictive results and can clarify the principles behind these predictions. Additionally, the obtained factors significantly help to dynamically present market changes over time in a format that is accessible and understandable to humans.

The main contributions of this study are summarized as follows:

- **New Task - Factor Extraction:** We introduce a novel task, factor extraction, aimed at extracting significant factors from textual data to assist in predicting time-series data. This task surpasses traditional keyphrase-based and sentiment-based methods in three key areas: the explainability of news content, the readability for humans, and the ability to explain stock price movements.
- **New Strategy - SKGP:** We propose Sequential Knowledge-Guided Prompting (SKGP) as an innovative strategy for leveraging background knowledge in the process of stock movement prediction. Unlike basic prompting methods, SKGP employs a fill-in-the-blank approach, utilizing minimal background knowledge to enhance the

richness of the prompt template.

- **New Framework - LLMFactor:** The LLMFactor framework, developed in this study, utilizes factors derived from LLMs to elucidate the dynamic temporal changes. This framework has proven to be a valuable tool in financial applications, providing deep insights into market trends and facilitating factor analysis.

2 Related Work

2.1 Stock Movement Prediction using Textual Data

With the advancement of natural language processing (NLP) techniques, many researchers leverage textual data to forecast stock market trends. (Xu and Cohen, 2018) uses tweets and historical prices to make temporally dependent predictions from stock data. (Luo et al., 2023) models the multimodality between financial text data and causality-enhanced stock correlations. These efforts aim to develop sophisticated models to enhance the prediction accuracy and consider the textual data as a whole.

Furthermore, certain studies focus on extracting more granular insights from textual data. (Zhou et al., 2021) identifies corporate events as the driving force behind stock movements while (Wan et al., 2021) demonstrates the significant association between strong media sentiment and market return. (Wang et al., 2020) enhances stock movement prediction by including expert opinions aggregated from various sources. These insights offer greater clarity to the understanding of market dynamics.

2.2 Time-series Forecasting with LLMs

LLMs, such as variants of the Generative Pre-trained Transformer (GPT), are constructed from wide-ranging and diverse datasets, significantly enriching their extensive knowledge base (Brown et al., 2020). Yet, their foundational structure, grounded in the transformer architecture, is not intrinsically designed for analyzing time-series data (Gruber et al., 2023). This limitation has sparked increased interest among scholars in exploring methods to tailor LLMs for proficient time-series prediction.

(Xue and Salim, 2023) proposes a prompt-based approach to transform numerical inputs and outputs into textual prompts, framing the forecasting task as a sentence-to-sentence conversion. This method allows language models to be directly applied to

❖ Task Definition

Input: stock's date, stock's news, past stock price movements → *Output:* today's stock price movement

❖ Sequential Knowledge-Guided Prompting

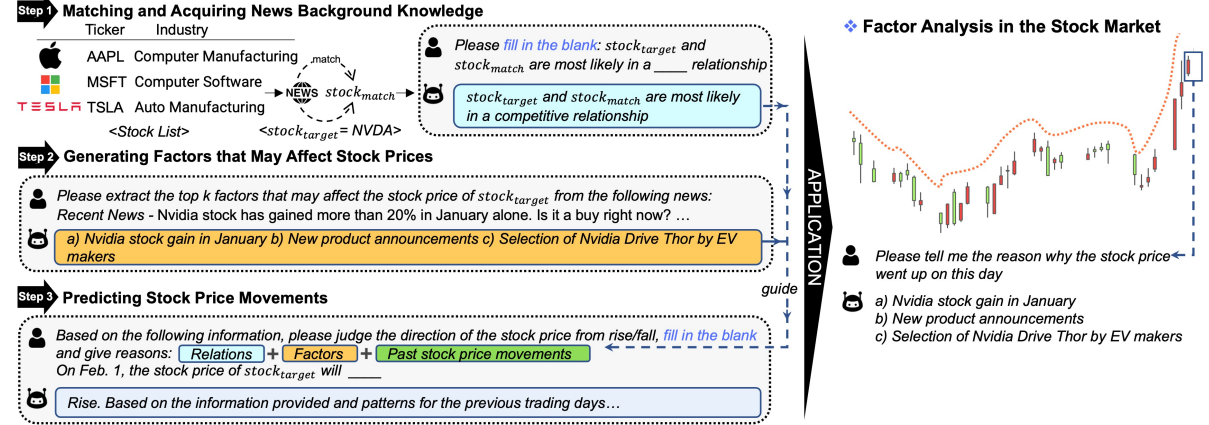


Figure 2: LLMFactor: From Sequential Knowledge-Guided Prompting (SKGP) to factor analysis in the stock market. The SKGP technique comprises three steps that sequentially generate knowledge and predict stock movement, guided by this knowledge. In applying factor analysis to stock movements, factors offer timely and concise explanations for their changes.

forecasting tasks. For financial time series forecasting, (Yu et al., 2023) employs prompts to generate summaries and keyphrases from LLMs, incorporating multiple data sources to enhance time-series forecasting. Their approaches offer valuable insights into stock market trends. However, their prompts contain excessive information, resulting in responses from LLMs that are relatively lacking in detail.

2.3 Prompt Engineering

Prompt engineering allows LLMs to address a variety of tasks effectively through carefully crafted prompts. However, identifying the optimal prompt for each specific task is a challenge (Liu et al., 2023). This challenge leads to the exploration of diverse prompting strategies. For instance, (Wei et al., 2022) introduces a chain of thought (CoT) approach, which involves a sequence of intermediate reasoning steps to enhance the capability of LLMs in complex reasoning tasks. Meanwhile, (Liu et al., 2022) proposes a generated knowledge prompting technique, which involves generating knowledge through a language model and then using this generated knowledge as supplementary input for answering questions. Another popular technique, known as Retrieval-Augmented Generation (RAG) (Lewis et al., 2020), incorporates external knowledge into LLMs through retrieve-augment-generate flows. Inspired by these techniques, we propose a novel Sequential Knowledge-Guided Prompting, designed

to enhance the accuracy of financial forecasts.

3 LLMFactor

The proposed LLMFactor is a comprehensive framework to predict and explain stock market trends as shown in Figure 2.

3.1 Task Definition

For a given stock, denoted as $stock_{target}$, we consider its associated news published on the target prediction date $date_{target}$, referred to as $news_{target}$, and its historical stock price sequence $P = \{P_1, P_2, \dots, P_t\}$, where t represents the window size. The task of predicting stock movement is formulated as a binary classification problem, where the stock price sequence is transformed into a series of stock movements, $\hat{P} = \{\hat{P}_1, \hat{P}_2, \dots, \hat{P}_t\}$. In this series, $\hat{P}_i = 1$ indicates a rise in the stock price from P_{i-1} to P_i , whereas $\hat{P}_i = 0$ indicates a fall. Our objective is to predict \hat{P}_{t+1} given $date_{target}$, $news_{target}$, and \hat{P} .

3.2 Sequential Knowledge-Guided Prompting

3.2.1 Matching and Acquiring News Background Knowledge

The foundation of our approach is the Sequential Knowledge-Guided Prompting (SKGP) strategy, which comprises three primary stages. Consider Nvidia (NVDA) as our example $stock_{target}$, the initial phase involves matching the stock with relevant news and acquiring background knowledge.

Let $S = \{(C_i, T_i, I_i) \mid i \in N, i \leq n\}$ be the stock list, where each tuple (C_i, T_i, I_i) consists of a company C_i , its ticker symbol T_i , and its industry I_i . These tuples are collected from the NYSE and NASDAQ exchanges, with n being a positive integer representing the total number of such tuples. We match S with $news_{target}$ to obtain $stock_{match} = \{S \cap news_{target}\}$. Following this, we prompt LLMs to obtain relations between $stock_{target}$ and $stock_{match}$. The prompting method is defined as $LLM : String \rightarrow String$, where the input is a *RelationTemplate* “Please fill in the blank: $stock_{target}$ and $stock_{match}$ are most likely in a ___ relationship,” and the output is the type of relationship, denoted as *relation*.

This method of deriving background knowledge about $news_{target}$ significantly improves our understanding of the news content. For example, if $stock_{target}$ and $stock_{match}$ are identified as competitors, information about $stock_{match}$ could potentially have a negative impact on $stock_{target}$. Furthermore, numerous studies have underscored the critical role of company relationships in predicting stock market movements (Hsu et al., 2021; Ang and Lim, 2022). Therefore, our fill-in-the-blank technique aims to constrain the response format, facilitating a direct and unambiguous identification of the relationship.

3.2.2 Generating Factors that May Affect Stock Prices

The next step for SKGP involves generating factors from $news_{target}$. The importance of these factors is threefold: 1) They are more closely associated with stock movements than keyphrases, sentiments, news summaries, or entire news articles, thus providing a higher likelihood of profitable market trend predictions. 2) Factors derived from news texts offer more immediate and detailed insights into stock price fluctuations compared to those obtained from other sources. 3) They improve the explainability of stock price trends and the rationale behind LLMs’ predictions.

To generate reliable factors, we instruct LLMs to analyze the news content and identify factors that could influence the stock price. This approach fully utilizes the intrinsic knowledge of LLMs. Considering the top k factors we aim to extract for the given $stock_{target}$ and $news_{target}$, the prompting method is described as $LLM(FactorTemplate) = factor$, where *FactorTemplate* is a structured sentence: “Please

extract the top k factors that may affect the stock price of $stock_{target}$ from the following news,” followed by $news_{target}$, and the output is the factors generated by the LLMs. The factors produced by LLMs are not restricted to the words found in the news; instead, LLMs consider the content of the news and its potential impact on stock movement, often summarizing the significant elements in the content. For example, the factor “Nvidia stock gain in January” mentioned in Figure 2 suggests that the event was Nvidia’s past stock performance, and the impact was a gain, which could positively influence Nvidia’s future stock price.

3.2.3 Predicting Stock Price Movements

To predict stock movement, we integrate news background knowledge and factors to guide LLMs. Meanwhile, we transform time-series data into a textual format for LLMs to understand. In a sequence of stock movements denoted by $\hat{P} = \{\hat{P}_1, \hat{P}_2, \dots, \hat{P}_t\}$, we introduce a function where $f(\hat{P}_i = 0)$ is assigned the value “fell” and $f(\hat{P}_i = 1)$ is assigned “rose,” thereby converting \hat{P} into a sequence of outcomes $TextMovement = \{f(\hat{P}_i) \mid i \in \{1, 2, \dots, t\}, f(\hat{P}_i) = \text{“fell” if } \hat{P}_i = 0 \text{ and “rose” if } \hat{P}_i = 1\}$. Given the textual stock movement sequence *TextMovement* and its date series $date = \{date_1, date_2, \dots, date_t\}$, the past stock price movements are transformed into a *TimeTemplate*, structured as “On $date_i$, the stock price of $stock_{target}$ $f(\hat{P}_i)$.”

Subsequently, we construct a *PriceTemplate* that includes an initial instruction, “Based on the following information, please judge the direction of the stock price as rise or fall, fill in the blank and give reasons,” followed by a concluding instruction, “On $date_i$, the stock price of $stock_{target}$ will ___.” By integrating the *relation*, *factor*, *TimeTemplate*, and *PriceTemplate*, we articulate the prompting method as $LLM(relation, factor, TimeTemplate, PriceTemplate) = prediction$. The *prediction* outcome specifies whether the stock price will “rise” or “fall,” along with the rationale for this inference. More details of templates are in Appendix A.

3.3 Factor Analysis in the Stock Market

The SKGP presents a powerful technique for predicting stock movements, and the factors derived from SKGP offer additional insights into stock market trends. For instance, as illustrated in Figure 2,

Statistics	StockNet	CMIN-US	CMIN-CN	EDT
Data Type	time series & text	time series & text	time series & text	text
Data Resource	price sequence & tweets	price sequence & tweets	price sequence & tweets	prices & news articles
Data Size	19,318	83,553	198,781	54,080
Stock Market	US	US	CN	US
Stock Number	87	110	300	4,228
Date Range	2014-01-01 to 2016-01-01	2018-01-01 to 2021-12-31	2018-01-01 to 2021-12-31	2020-03-01 to 2021-05-06

Table 1: Statistics of four benchmark datasets: StockNet, CMIN-US, and CMIN-CN show consistent processing across each dataset, with every stock represented by time-series data that includes both tweets and stock prices. In contrast, the EDT dataset is composed of news articles, with each article linked to the relevant stock and its corresponding stock price.

factor analysis can be applied to the stock market. Taking Nvidia’s stock price trend as an example, after a consistent increase over the past five days, the day highlighted with a blue box also exhibits an upward movement. To explain this occurrence, our LLMFactor identifies a concise set of factors, such as “Nvidia stock gain in January, new product announcements, and selection of Nvidia Drive Thor by EV makers.”

4 Experiments

By conducting experiments on benchmark datasets, we aim to address the following research question: Does LLMFactor offer improved prediction accuracy and more insightful explanations for stock market movements?

4.1 Datasets

We evaluate LLMFactor on four benchmark datasets, in alignment with the current SOTAs: a) **StockNet** (Xu and Cohen, 2018) includes 87 stocks from 9 industries, accompanied by stock-related tweets and historical price data from 2014-01-01 to 2016-01-01, in the US stock market; b) **CMIN-US** (Luo et al., 2023) comprises the top 110 stocks along with their tweets and historical price data from 2018-01-01 to 2021-12-31, in the US stock market; c) **CMIN-CN** (Luo et al., 2023) consists of the 300 stocks in the CSI300 index ¹, with their tweets and historical price data from 2018-01-01 to 2021-12-31, in the Chinese stock market; d) **EDT** (Zhou et al., 2021) includes 54,080 news articles from 2020-03-01 to 2021-05-06, along with related stock and stock price information in the US stock market. Unlike StockNet, CMIN-US, and CMIN-CN, which focus on time-series forecasting, the EDT dataset is centred on news content. The

¹A market capitalization-weighted index crafted to mirror the performance of the leading 300 stocks traded on the Shanghai Stock Exchange and the Shenzhen Stock Exchange.

statistics of all datasets are summarized in Table 1.

4.2 Evaluation Metrics

Following previous studies by (Xu and Cohen, 2018; Luo et al., 2023; Zhou et al., 2021), we apply Accuracy (ACC) and Matthews Correlation Coefficient (MCC) as the evaluation metrics. Given the confusion matrix $\begin{bmatrix} tp & fp \\ fn & tn \end{bmatrix}$, ACC and MCC are defined as:

$$ACC = \frac{tp + tn}{tp + tn + fp + fn} \quad (1)$$

$$MCC = \frac{tp \times tn - fp \times fn}{\sqrt{(tp + fp)(tp + fn)(tn + fp)(tn + fn)}} \quad (2)$$

4.3 Baselines

Keyphrase-based. We employ several models to identify keyphrases in text for prediction purposes:

- **PromptRank**(Kong et al., 2023): An unsupervised approach that utilizes pre-trained language models (PLMs) with a focus on leveraging prompt-based techniques for keyphrase extraction.
- **KeyBERT**(Grootendorst, 2020): Leverages BERT embeddings to identify key sub-phrases in documents through cosine similarity, emphasizing the most representative phrases.
- **YAKE**(Campos et al., 2020): A lightweight unsupervised method for keyword extraction, employing statistical features from individual documents to pinpoint the most pertinent keywords.
- **TextRank**(Mihalcea and Tarau, 2004): An unsupervised approach identifies keyphrases based on the co-occurrence graph of words in a text.
- **TopicRank**(Bougouin et al., 2013): Groups similar expressions into topics and ranks these topics to identify keyphrases.
- **SingleRank**(Wan and Xiao, 2008): An extension of TextRank that weights edges in the co-occurrence graph based on the number of co-occurrences.

- **TFIDF**([Sparck Jones, 1972](#)): A traditional unsupervised approach that evaluates the importance of a word within a document in a corpus.

Sentiment-based. We employ various models to analyze sentiment in texts and predict stock market movements:

- **EDT**([Zhou et al., 2021](#)): Utilizes a bi-level event detection framework based on BERT to classify corporate events impacting stock prices.
- **FinGPT**([Zhang et al., 2023](#)): Employs instruction tuning on LLMs to enhance financial sentiment analysis capabilities.
- **GPT-4-turbo**: Builds upon GPT-4’s capabilities, enhancing efficiency and performance, especially for processing longer texts.
- **GPT-4**: Represents a significant leap forward from previous iterations, with improved processing speeds and enhanced handling of extended text inputs.
- **GPT-3.5-turbo**: Known for its comprehensive understanding and human-like text generation, with optimized processing for multiple tasks.
- **RoBERTa**: Applies Financial RoBERTa² for sentiment analysis in English texts and a fine-tuned sentiment-specific RoBERTa model([Wang et al., 2022](#)) for Chinese texts.
- **FinBERT**: A BERT model fine-tuned for financial sentiment classification, available in both English ([Araci, 2019](#)) and Chinese versions³.

Time-based. We employ models that leverage both textual and time-series data to predict stock movement:

- **CMIN**([Luo et al., 2023](#)): An end-to-end deep neural network that models the multimodality between financial text data and causality-enhanced stock correlations.
- **StockNet**([Xu and Cohen, 2018](#)): A deep generative model that jointly exploits textual and price signals for stock prediction.

4.4 Implementation Details

For keyphrase-based methods, the datasets are divided into two subsets based on labels: one containing data labeled with 0 (indicating a fall) and the other with data labeled with 1 (indicating a rise). Keyphrase extraction is applied to each subset to

²<https://huggingface.co/soleimanian/financial-roberta-large-sentiment>

³<https://huggingface.co/bardsai/finance-sentiment-zh-base>

identify negative and positive keyphrases for each stock. The score for a given text is calculated as follows: Let POS be the set of positive keyphrases and NEG be the set of negative keyphrases. For a piece of text T , the score S is computed as:

$$S = \sigma \left(\sum_{p \in POS} I(p \in T) - \sum_{n \in NEG} I(n \in T) \right) \quad (3)$$

where $\sigma(x)$ is the sigmoid function defined as $\sigma(x) = \frac{1}{1+e^{-x}}$, and $I(\cdot)$ is an indicator function that equals 1 if the condition is true and 0 otherwise. The condition $p \in T$ (or $n \in T$) denotes the presence of a positive (or negative) keyphrase in the text T . For sentiment analysis, we predict the sentiment of the text as either positive or negative and anticipate a rise if the text is positive and a fall if it is negative, in alignment with baseline settings. However, since the EDT dataset contains only textual data, we do not employ time-based methods on it. Additionally, for our factor-based methods, we execute the entire SKGP process on benchmark datasets but exclude the use of the *TimeTemplate* when assessing the EDT dataset. For a fair comparison, we ensure consistency in the processing of baselines according to their original settings. We employ gpt-3.5-turbo-1106, gpt-4, and gpt-4-1106-preview through the API⁴ for factor-based methods. The parameters are set as follows: 1) the window size t is 5; 2) the number of keyphrases and factors k is 5; 3) the batch size is 64 for BERT series models and 5 for GPT series models. Our experiments are conducted with an NVIDIA RTX A6000 GPU. The details of the pre-trained models used for the baselines are summarized in Appendix B.

4.5 Results

Table 2 presents the results for LLMFactor and various baselines across four datasets. Among all evaluated approaches, our factor-based methods surpass the others, with time-based, sentiment-based, and keyphrase-based methods following in performance. Notably, LLMFactor achieves superior performance over the SOTAs, with improvements of **2.9%**, **0.4%**, **11%**, and **4.8%** in MCC across the four datasets, respectively.

The minimal performance differences among keyphrase-based models suggest their limited effectiveness in forecasting stock movements using keyphrases, as keyphrases primarily summarize the main content of the text rather than provide actionable insights on stock prices. Addition-

⁴<https://platform.openai.com/docs/models>

Method	Model	StockNet		CMIN-US		CMIN-CN		EDT	
		ACC.	MCC.	ACC.	MCC.	ACC.	MCC.	ACC.	MCC.
Keyphrase-based	Promptrank (ACL2023)	51.24	0.010	53.28	0.001	50.21	-0.003	51.78	0.014
	KeyBERT	51.95	0.012	53.40	0.009	50.23	-0.004	51.84	0.002
	YAKE	51.91	0.005	53.13	0.001	50.20	0.001	51.88	0.004
	TextRank	51.00	0.003	53.99	0.060	50.38	0.006	51.76	0.003
	TopicRank	51.92	0.008	53.75	0.034	50.26	-0.002	51.80	0.000
	SingleRank	50.32	0.005	53.33	0.004	50.29	0.002	51.85	0.004
	TFIDF	51.86	0.001	53.71	0.018	50.27	-0.002	51.86	0.017
Sentiment-based	EDT (ACL2021)	40.31	-0.066	49.86	-0.004	40.00	0.021	75.67	0.026
	FinGPT	54.91	0.083	59.98	0.182	55.78	0.120	53.86	0.035
	GPT-4-turbo	53.56	0.060	64.61	0.284	56.94	0.135	55.37	0.057
	GPT-4	53.88	0.062	62.18	0.260	56.96	0.136	50.94	0.031
	GPT-3.5-turbo	52.31	0.044	56.10	0.156	56.68	0.124	54.34	0.040
	RoBERTa	54.46	0.088	57.75	0.138	52.24	0.064	53.66	0.029
	FinBERT	55.42	0.111	58.26	0.158	55.98	0.121	54.98	0.043
Time-based	CMIN (ACL 2023)	62.69	0.209	53.43	0.046	55.28	0.111	-	-
	StockNet (ACL 2018)	58.23	0.081	52.46	0.022	54.53	0.045	-	-
Factor-based(ours)	LLMFactor _{GPT-4-turbo}	<u>65.81</u>	<u>0.228</u>	61.71	0.228	60.59	0.245	59.09	<u>0.082</u>
	LLMFactor _{GPT-4}	66.32	0.238	<u>65.26</u>	<u>0.284</u>	<u>57.16</u>	<u>0.196</u>	<u>60.83</u>	0.105
	LLMFactor _{GPT-3.5-turbo}	57.59	0.145	66.42	0.288	56.11	0.139	58.11	0.097

Table 2: Results for LLMFactor and other baselines. The ACC. is shown in percentages (%). Bold text indicates the best results, while underlined text indicates the second-best. We omit the comparison of time-based methods for the EDT dataset due to its lack of time-series data.

Method	Model	StockNet		CMIN-US		CMIN-CN		EDT	
		ACC.	MCC.	ACC.	MCC.	ACC.	MCC.	ACC.	MCC.
LLMFactor	Price	52.16	0.041	55.59	0.135	51.76	0.048	-	-
	+Factor	58.04	0.166	61.68	0.241	55.71	0.119	55.93	0.077
	+Factor+Relation	63.24	0.203	64.46	0.267	57.96	0.194	59.35	0.095

Table 3: The ablation study results for LLMFactor, summarizing average values across $LLMFactor_{GPT-4-turbo}$, $LLMFactor_{GPT-4}$, and $LLMFactor_{GPT-3.5-turbo}$. The ACC. is expressed in percentages (%), with the highest performance bolded for emphasis. The analysis excludes price data comparisons in the EDT dataset due to the absence of time-series data.

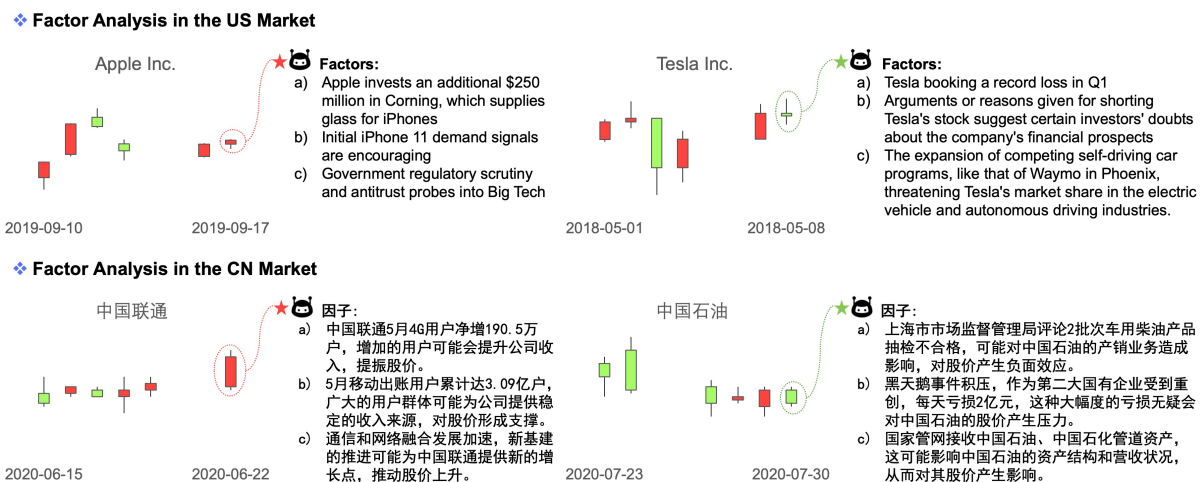


Figure 3: Analyzing the impact of LLMFactor on US and CN markets: LLMFactor is capable of identifying key factors that influence daily fluctuations in stock prices for individual stocks.

ally, sentiment-based methods demonstrate varying performance across different models. LLMs like GPT-4 and GPT-4-turbo excel in identifying sentiment within the text. Although the EDT model achieves the highest accuracy on its specific dataset, its low MCC indicates imbalanced model performance. Unlike sentiment-based methods geared towards efficiently extracting relevant sentiment, time-based methods focus on integrating multiple data types. Since time-based methods consider the text in its entirety without detailed analysis, their performance is on par with that of sentiment-based methods, which may lack temporal context but include detailed textual information. However, our LLMFactor model stands out among other SOTAs due to the SKGP technique, which not only identifies significant factors influencing stock prices from textual data but also incorporates relation and temporal information, thereby enhancing its ability to filter out irrelevant content and provide a more comprehensive analysis.

Analyzing average scores from diverse datasets with the use of three factor-based models, our LLMFactor demonstrates an average accuracy of over 63% and an average MCC of over 0.2 for both the StockNet and CMIN-US datasets, tailored to the US market. However, its performance slightly declines on the CMIN-CN dataset, focused on the CN market, with an average accuracy of 58% and an MCC of 0.19. This variation indicates that the GPT series models are potentially more proficient in processing English-language text. The EDT dataset, which consists solely of stock-related news without historical price data, poses additional challenges, resulting in a marginally lower average accuracy of 59% and an MCC of 0.1 for the LLMFactor. The absence of historical price information in the EDT dataset likely diminishes the LLMFactor’s effectiveness, underscoring the importance of comprehensive data for financial market analysis.

4.6 Data Analysis

4.6.1 Ablation Study

We conduct an ablation study on the LLMFactor to demonstrate the improvement contributed by each layer. The first layer utilizes only the *TimeTemplate*, which incorporates past price movements. The second layer incorporates additional *factors* extracted, and the third layer further includes *relations* obtained. We average the scores for ACC and MCC from $LLM_{GPT-4-turbo}$,

LLM_{GPT-4} , and $LLM_{GPT-3.5-turbo}$, as shown in Table 3. The price layer accounts for approximately 86% and 32% of the total performance in ACC and MCC, respectively, while the factor layer contributes an improvement of 9% in ACC and 46% in MCC. The relation layer leads to a 5% and 22% enhancement in ACC and MCC, respectively. These results imply that the factor layer contributes most significantly to the overall performance of the proposed LLMFactor.

We also conduct experiments on different types of *FactorTemplate*, as shown in Appendix C.

4.6.2 Case Study

In this section, we highlight the practical effectiveness of LLMFactor through an illustrative application. Figure 3 presents a case study where we apply factor analysis to selected stocks in the US and CN markets. For instance, on September 17, 2019, Apple Inc.’s stock (AAPL) experienced an increase, attributed to significant developments such as a \$250 million investment in its supplier, Corning Incorporated, and the positive initial demand for the iPhone 11. Conversely, on May 8, 2018, Tesla Inc.’s stock (TSLA) experienced a decline, mainly due to a record quarterly loss, growing negative concerns among investors, and the increasing competition in the autonomous and electric vehicle markets. This case study demonstrates LLMFactor’s ability to effectively integrate background knowledge of company affiliations with both historical news and price data. By analyzing these diverse data sources, LLMFactor significantly enhances the interpretability of stock market dynamics.

5 Conclusion

In this study, we introduce LLMFactor, an innovative framework designed to predict and elucidate stock market trends accurately. At the core of LLMFactor is the Sequential Knowledge-Guided Prompting (SKGP) strategy, which integrates background knowledge, stock-related factors, and temporal data to forecast stock movements. Through rigorous testing on four benchmark datasets, LLMFactor has proven its superiority over SOTAs that rely on keyphrases, sentiment analysis, and multi-modal data inputs. The application of factor analysis highlights LLMFactor’s novelty and effectiveness, establishing it as a powerful tool for financial analysis. This research represents a significant advancement in leveraging LLMs for transparent and explainable financial forecasting.

6 Limitations

We recognize three main limitations:

- Although LLMFactor’s foundation relies on extensively discussed factors, converting temporal data into text format remains vital for financial forecasting. Future studies could explore additional methods for this conversion and their integration into LLMs.
- Due to the variable nature of LLM responses, replicating our experimental results precisely is challenging. In future research, we aim to explore further methodologies to improve reproducibility.
- Evaluating LLMFactor against benchmark datasets enables comparison with leading models, but it’s essential to carefully assess the quality of factors extracted from texts of various lengths and types. This research marks the beginning of a novel approach to applying factors in finance. Moving forward, we will concentrate on enhancing the quality of these factors.

Acknowledgements

We sincerely thank the anonymous reviewers for their reviews. This work was supported by JST SPRING, Grant Number JPMJSP2108, and JST-PRESTO, Grant Number JPMJPR2267.

References

- Noujoud Ahbali, Xinyuan Liu, Albert Nanda, Jamie Stark, Ashit Talukder, and Rupinder Paul Khandpur. 2022. [Identifying Corporate Credit Risk Sentiments from Financial News](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Track*, pages 362–370, Hybrid: Seattle, Washington + Online. Association for Computational Linguistics.
- Gary Ang and Ee-Peng Lim. 2022. [Guided attention multimodal multitask financial forecasting with inter-company relationships and global and local news](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6313–6326, Dublin, Ireland. Association for Computational Linguistics.
- Dogu Araci. 2019. [Finbert: Financial sentiment analysis with pre-trained language models](#). *arXiv preprint arXiv:1908.10063*.
- Astrid Bertrand, James R Eagan, and Winston Maxwell. 2023. [Questioning the ability of feature-based explanations to empower non-experts in robo-advised financial decision-making](#). In *Proceedings of the*

2023 ACM Conference on Fairness, Accountability, and Transparency, pages 943–958.

- Florian Boudin. 2016. [pke: an open source python-based keyphrase extraction toolkit](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: System Demonstrations*, pages 69–73, Osaka, Japan. The COLING 2016 Organizing Committee.
- Adrien Bougouin, Florian Boudin, and Béatrice Daille. 2013. [TopicRank: Graph-based topic ranking for keyphrase extraction](#). In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 543–551, Nagoya, Japan. Asian Federation of Natural Language Processing.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. [Language models are few-shot learners](#). *Advances in neural information processing systems*, 33:1877–1901.
- Ricardo Campos, Vítor Mangaravite, Arian Pasquali, Alípio Jorge, Célia Nunes, and Adam Jatowt. 2020. [Yake! keyword extraction from single documents using multiple local features](#). *Information Sciences*, 509:257–289.
- Eugene F Fama. 1970. Efficient capital markets: A review of theory and empirical work. *The journal of Finance*, 25(2):383–417.
- Fuli Feng, Xiangnan He, Xiang Wang, Cheng Luo, Yiqun Liu, and Tat-Seng Chua. 2019. [Temporal relational ranking for stock prediction](#). *ACM Transactions on Information Systems (TOIS)*, 37(2):1–30.
- Maarten Grootendorst. 2020. [Keybert: Minimal keyword extraction with bert](#).
- Nate Gruver, Marc Anton Finzi, Shikai Qiu, and Andrew Gordon Wilson. 2023. [Large language models are zero-shot time series forecasters](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Pei-Yi Hao, Chien-Feng Kung, Chun-Yang Chang, and Jen-Bing Ou. 2021. [Predicting stock price trends based on financial news articles and using a novel twin support vector machine with fuzzy hyperplane](#). *Applied Soft Computing*, 98:106806.
- Bruno Miranda Henrique, Vinicius Amorim Sobreiro, and Herbert Kimura. 2019. [Literature review: Machine learning techniques applied to financial market prediction](#). *Expert Systems with Applications*, 124:226–251.
- Yi-Ling Hsu, Yu-Che Tsai, and Cheng-Te Li. 2021. [Fingat: Financial graph attention networks for recommending top-k profitable stocks](#). *IEEE Transactions on Knowledge and Data Engineering*, 35(1):469–481.

- Aobo Kong, Shiwan Zhao, Hao Chen, Qicheng Li, Yong Qin, Ruiqi Sun, and Xiaoyan Bai. 2023. [PromptRank: Unsupervised keyphrase extraction using prompt](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9788–9801, Toronto, Canada. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Wei Li, Ruihan Bao, Keiko Harimoto, Deli Chen, Jingjing Xu, and Qi Su. 2021. [Modeling the stock relation with graph network for overnight stock movement prediction](#). In *Proceedings of the twenty-ninth international conference on international joint conferences on artificial intelligence*, pages 4541–4547.
- Jiacheng Liu, Alisa Liu, Ximing Lu, Sean Welleck, Peter West, Ronan Le Bras, Yejin Choi, and Hannaneh Hajishirzi. 2022. [Generated knowledge prompting for commonsense reasoning](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3154–3169, Dublin, Ireland. Association for Computational Linguistics.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#). *ACM Computing Surveys*, 55(9):1–35.
- Andrew W Lo. 2004. The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *Journal of Portfolio Management*, *Forthcoming*.
- Di Luo, Weiheng Liao, Shuqi Li, Xin Cheng, and Rui Yan. 2023. [Causality-guided multi-memory interaction network for multivariate stock price movement prediction](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12164–12176, Toronto, Canada. Association for Computational Linguistics.
- Rada Mihalcea and Paul Tarau. 2004. [TextRank: Bringing order into text](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 404–411, Barcelona, Spain. Association for Computational Linguistics.
- Thien Hai Nguyen, Kiyooki Shirai, and Julien Velcin. 2015. [Sentiment analysis on social media for stock movement prediction](#). *Expert Systems with Applications*, 42(24):9603–9611.
- Jigar Patel, Sahil Shah, Priyank Thakkar, and Ketan Kotecha. 2015. [Predicting stock and stock price index movement using trend deterministic data preparation and machine learning techniques](#). *Expert systems with applications*, 42(1):259–268.
- Karen Sparck Jones. 1972. A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation*, 28(1):11–21.
- Allan Timmermann and Clive WJ Granger. 2004. Efficient market hypothesis and forecasting. *International Journal of forecasting*, 20(1):15–27.
- Manuel R Vargas, Carlos EM Dos Anjos, Gustavo LG Bichara, and Alexandre G Evsukoff. 2018. [Deep learning for stock market prediction using technical indicators and financial news articles](#). In *2018 international joint conference on neural networks (IJCNN)*, pages 1–8. IEEE.
- Xiaojun Wan and Jianguo Xiao. 2008. [CollabRank: Towards a collaborative approach to single-document keyphrase extraction](#). In *Proceedings of the 22nd International Conference on Computational Linguistics (Coling 2008)*, pages 969–976, Manchester, UK. Coling 2008 Organizing Committee.
- Xingchen Wan, Jie Yang, Slavi Marinov, Jan-Peter Caließ, Stefan Zohren, and Xiaowen Dong. 2021. [Sentiment correlation in financial news networks and associated market movements](#). *Scientific reports*, 11(1):3062.
- Heyuan Wang, Tengjiao Wang, and Yi Li. 2020. [Incorporating expert-based investment opinion signals in stock prediction: A deep learning framework](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 971–978.
- Junjie Wang, Yuxiang Zhang, Lin Zhang, Ping Yang, Xinyu Gao, Ziwei Wu, Xiaoqun Dong, Junqing He, Jianheng Zhuo, Qi Yang, Yongfeng Huang, Xiayu Li, Yan-Ze Wu, Junyu Lu, Xinyu Zhu, Weifeng Chen, Ting-Ting Han, Kunhao Pan, Rui Wang, Hao Wang, Xiaojun Wu, Zhong Zeng, Chong-An Chen, Ruyi Gan, and Jiaying Zhang. 2022. [Fengshenbang 1.0: Being the foundation of chinese cognitive intelligence](#). *ArXiv*, abs/2209.02970.
- Zhicheng Wang, Biwei Huang, Shikui Tu, Kun Zhang, and Lei Xu. 2021. [Deeptrader: a deep reinforcement learning approach for risk-return balanced portfolio management with market conditions embedding](#). In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 643–650.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Ming Xu. 2023. [pke_zh: Python keyphrase extraction toolkit for chinese](#).
- Yumo Xu and Shay B. Cohen. 2018. [Stock movement prediction from tweets and historical prices](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1970–1979, Melbourne, Australia. Association for Computational Linguistics.

Hao Xue and Flora D Salim. 2023. [Promptcast: A new prompt-based learning paradigm for time series forecasting](#). *IEEE Transactions on Knowledge and Data Engineering*.

Jingming Xue, En Zhu, Qiang Liu, and Jianping Yin. 2018. [Group recommendation based on financial social network for robo-advisor](#). *IEEE Access*, 6:54527–54535.

Xinli Yu, Zheng Chen, and Yanbin Lu. 2023. [Harnessing LLMs for temporal data - a study on explainable financial time series forecasting](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 739–753, Singapore. Association for Computational Linguistics.

Boyuan Zhang, Hongyang Yang, and Xiao-Yang Liu. 2023. [Instruct-fingpt: Financial sentiment analysis by instruction tuning of general-purpose large language models](#). *FinLLM Symposium at IJCAI 2023*.

Zhihan Zhou, Liqian Ma, and Han Liu. 2021. [Trade the event: Corporate events detection for news-based event-driven trading](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2114–2124, Online. Association for Computational Linguistics.

A Templates of SKGP

The fundamental templates employed at each step of SKGP are detailed in Table 5 for English and Table 6 for Chinese. In the final step of SKGP, we compile the responses from previous steps to formulate prompts for LLMs. Nonetheless, due to the variable nature of LLM responses, the exact reproducibility of results cannot be guaranteed. Instead, we demonstrate the overall superior performance of the LLMFactor.

B Details of Pre-trained Models

We employ the pke (Boudin, 2016) and pke_zh (Xu, 2023) toolkits to access keyphrase-based methodologies, ensuring alignment with established baselines. Additionally, configuration details for sentiment-based models are outlined in Table 4, provided by HuggingFace. The models specified at the top are designated for English texts, whereas the ones at the bottom cater to Chinese texts.

Model	HuggingFace Key	Model Size
FinBERT	ProsusAI/finbert	438MB
	bardsai/finance-sentiment-zh-base	409MB
RoBERTa	soleimanian/financial-roberta-large-sentiment	1.42GB
	IDEA-CCNL/Erlangshen-Roberta-110M-Sentiment	409MB
FinGPT	FinGPT/fingpt-sentiment_llama2-13b_lora	14.36GB
	oliverwang15/FinGPT_ChatGLM2_Sentiment_Instruction_LoRA_FT	7.88GB

Table 4: Details of pre-trained models

C Comparative Experimental Results of Factor Templates

The experimental results for various factor templates, utilizing three factor-based models: $LLMFactor_{GPT-4-turbo}$, $LLMFactor_{GPT-4}$, and $LLMFactor_{GPT-3.5-turbo}$, are showcased in Table 7 for English templates and in Table 8 for Chinese templates. The datasets tested include CMIN-US and CMIN-CN. Within the English templates, the last template achieves the highest performance, while among the Chinese templates, the initial template emerges as the most effective. The consistent effectiveness of LLMFactor across different templates underscores its robustness in predicting stock market movements.

Step	EN-Template
Step1	Please fill in the blank and return a complete sentence: $stock_{target}$ and $stock_{match}$ are most likely in a ___ relationship.
Step2	Please extract the top k factors that may affect the stock price of $stock_{target}$ from the following news.
Step3	Based on the following information, please judge the direction of the stock price from rise/fall, fill in the blank and give reasons. These are the main factors that may affect this stock's price recently: $factor$. These are the connections between the companies that have appeared in the news: $relation$. On $date_{i-5}$, the stock price of $stock_{target}$ $f(P_{i-5})$. On $date_{i-4}$, the stock price of $stock_{target}$ $f(P_{i-4})$. On $date_{i-3}$, the stock price of $stock_{target}$ $f(P_{i-3})$. On $date_{i-2}$, the stock price of $stock_{target}$ $f(P_{i-2})$. On $date_{i-1}$, the stock price of $stock_{target}$ $f(P_{i-1})$. On $date_i$, the stock price of $stock_{target}$ will ___.

Table 5: SKGP Templates in English

Step	CN-Template
Step1	请填写并返回完整的句子: $stock_{target}$ 和 $stock_{match}$ 最可能是__关系。
Step2	请从以下新闻中提取可能影响 $stock_{target}$ 股价的前 k 个因素。
Step3	根据以下信息, 请判断股票价格是上涨还是下跌, 填写在空白处并给出理由。 这些是最近可能影响该股票价格的主要因素: $factor$ 这些是新闻中出现过的公司之间的关系: $relation$ 在 $date_{i-5}$, $stock_{target}$ 的股价 $f(P_{i-5})$ 。 在 $date_{i-4}$, $stock_{target}$ 的股价 $f(P_{i-4})$ 。 在 $date_{i-3}$, $stock_{target}$ 的股价 $f(P_{i-3})$ 。 在 $date_{i-2}$, $stock_{target}$ 的股价 $f(P_{i-2})$ 。 在 $date_{i-1}$, $stock_{target}$ 的股价 $f(P_{i-1})$ 。 在 $date_i$, $stock_{target}$ 的股价将___。

Table 6: SKGP Templates in Chinese

Factor Template	LLMFactor _{GPT-4-turbo}		LLMFactor _{GPT-4}		LLMFactor _{GPT-3.5-turbo}	
	ACC.	MCC.	ACC.	MCC.	ACC.	MCC.
Please extract the top k factors that may affect the stock price of $stock_{target}$ from the following news	61.71	0.228	65.26	0.284	66.42	0.288
Please identify the primary top k factors influencing $stock_{target}$'s stock price based on the news provided	66.98	0.292	64.83	0.246	65.26	0.298
Please analyze the provided news and pinpoint the top k major factors impacting the stock price of $stock_{target}$	65.29	0.295	69.21	0.312	66.56	0.293

Table 7: Experimental Results for Different Factor Templates in English

Factor Template	LLMFactor _{GPT-4-turbo}		LLMFactor _{GPT-4}		LLMFactor _{GPT-3.5-turbo}	
	ACC.	MCC.	ACC.	MCC.	ACC.	MCC.
请从以下新闻中提取可能影响 $stock_{target}$ 股价的前 k 个因素	60.59	0.245	57.16	0.196	56.11	0.139
根据提供的新闻, 请识别出影响 $stock_{target}$ 股价的主要 k 个因素	57.92	0.147	64.79	0.109	65.13	0.139
请分析所提供的新闻并找出影响 $stock_{target}$ 股价的前 k 个主要因素	64.29	0.160	59.33	0.053	59.49	0.033

Table 8: Experimental Results for Different Factor Templates in Chinese