

ICNLP 2025
The 7th International Conference on Natural Language Processing

ADAM HORN AND TIM SCHLIPPE

INVESTIGATING THE PREDICTIVE CAPABILITIES OF LARGE LANGUAGE MODELS IN DAY TRADING BY LEVERAGING MULTIMODAL DATA

Guangzhou, China

March 23, 2025

AGENDA

Introduction

1

Related Work

2

Experimental Setup

3

Experiments and Results

4

Conclusion and Future Work

5

1

INTRODUCTION

MOTIVATION: Importance of Financial Market Forecasting




Crucial for

- risk management
- trading strategies
- investment decisions

Image Sources: Image supplied by Microsoft 365. Text Sources: Araci (2019); Cao et al. (2024).

MOTIVATION: Challenges of Financial Market Forecasting



Challenging due to

- growing complexity
- volatility of global markets

Image Sources: Image supplied by Microsoft 365. Text Sources: Araci (2019); Cao et al. (2024).

MOTIVATION: Traditional Approaches



Image Sources: Image supplied by Microsoft 365. Text Sources: Kou et al. (2024).

e.g.

- Technical analysis
- Fundamental analysis

→ Fail to adapt
to market shifts

→ Struggle with
vast data integration

MOTIVATION: Large Language Models



e.g.

- Technical analysis
- Fundamental analysis

→ vast data integ

2

RELATED WORK

RELATED WORK

- Early forecasting relied on mathematical models (technical & fundamental analysis).
→ But struggled in dynamic markets (Cao et al., 2024).

RELATED WORK

- Early forecasting relied on mathematical models (technical & fundamental analysis).
➔ But struggled in dynamic markets (Cao et al., 2024).
- Transformer models (BERT, RoBERTa, FinBERT) improved text-based forecasting.
➔ But required extensive fine-tuning and computational resources.
(Araci, 2019; Cao et al., 2024; Guo & Hauptmann, 2024)

RELATED WORK

- Early forecasting relied on mathematical models (technical & fundamental analysis).
➔ But struggled in dynamic markets (Cao et al., 2024).
- Transformer models (BERT, RoBERTa, FinBERT) improved text-based forecasting.
➔ But required extensive fine-tuning and computational resources.
(Araci, 2019; Cao et al., 2024; Guo & Hauptmann, 2024)
- LLMs offer flexibility through prompting, handling unstructured data like news for sentiment analysis and stock forecasting (Cao et al., 2024, Daube & Krivenkov, 2024).

RELATED WORK

- Early forecasting relied on mathematical models (technical & fundamental analysis).
➔ But struggled in dynamic markets (Cao et al., 2024).
- Transformer models (BERT, RoBERTa, FinBERT) improved text-based forecasting.
➔ But required extensive fine-tuning and computational resources.
(Araci, 2019; Cao et al., 2024; Guo & Hauptmann, 2024)
- LLMs offer flexibility through prompting, handling unstructured data like news for sentiment analysis and stock forecasting (Cao et al., 2024, Daube & Krivenkov, 2024).
- Multi-agent systems with LLMs outperformed traditional models in financial prediction.
(Kou et al., 2024)

RELATED WORK

- Early forecasting relied on mathematical models (technical & fundamental analysis).
➔ But struggled in dynamic markets (Cao et al., 2024).
- Transformer models (BERT, RoBERTa, FinBERT) improved text-based forecasting.
➔ But required extensive fine-tuning and computational resources.
(Araci, 2019; Cao et al., 2024; Guo & Hauptmann, 2024)
- LLMs offer flexibility through prompting, handling unstructured data like news for sentiment analysis and stock forecasting (Cao et al., 2024, Daube & Krivenkov, 2024).
- Multi-agent systems with LLMs outperformed traditional models in financial prediction.
(Kou et al., 2024)
- Advanced approaches (e.g., “Denoising-then-Voting,” QLoRA) enhanced prediction accuracy.
(Daube & Krivenkov, 2024; Deng et al., 2024; Ni et al., 2024)

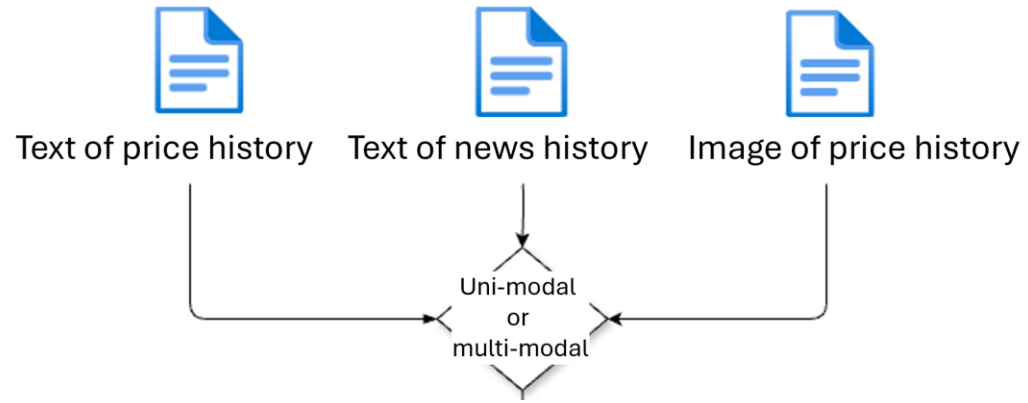
RELATED WORK

- Early forecasting relied on mathematical models (technical & fundamental analysis).
→ But struggled in dynamic markets (Cao et al., 2024).
- Traditional models (e.g., LSTM, RoBERTa, FinBERT) improved text-based forecasting.
→ But required significant data and computational resources.
- LLMs (e.g., GPT-4, Gemini) can process news for sentiment analysis and stock price forecasting.
- Multi-agent systems with LLMs outperformed traditional models (Kou et al., 2024).
- Advanced approaches (e.g., “Denoising-then-Voting,” QLoRA) enhanced LLM performance (Daube & Krivenkov, 2024; Deng et al., 2024; Ni et al., 2024).

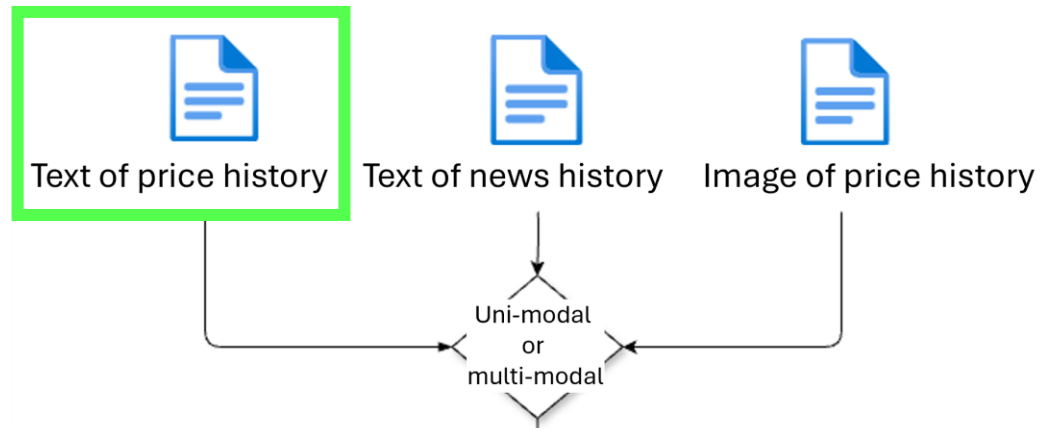
NO PRIOR STUDY COMPARES MULTIPLE STATE-OF-THE-ART LLMS FOR MULTIMODAL STOCK PRICE FORECASTING.

EXPERIMENTAL SETUP

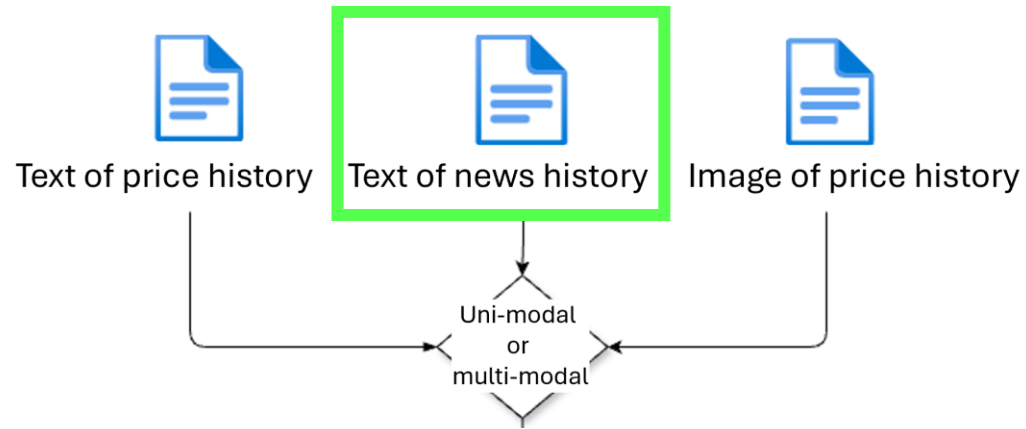
EXPERIMENTAL SETUP: OVERVIEW



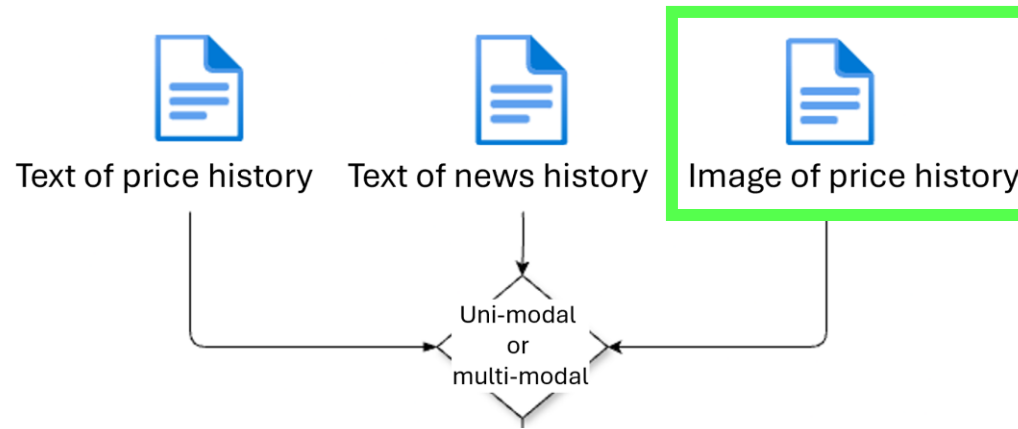
EXPERIMENTAL SETUP: OVERVIEW



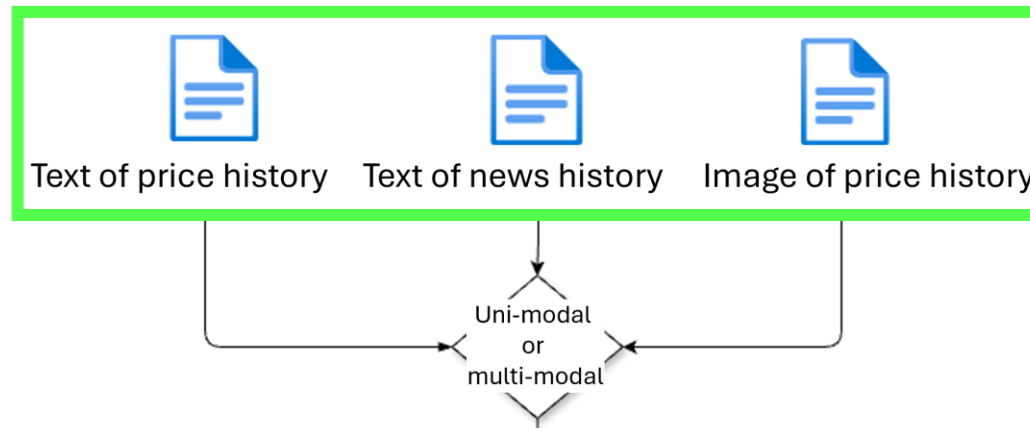
EXPERIMENTAL SETUP: OVERVIEW



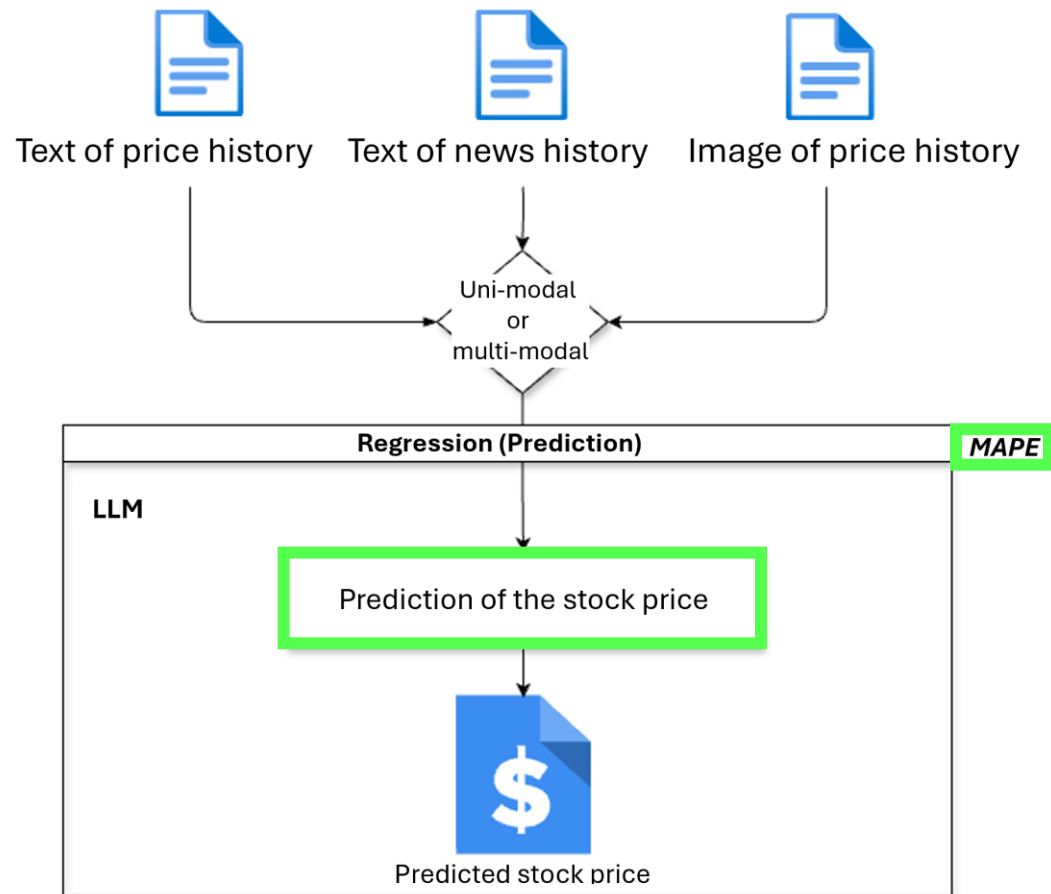
EXPERIMENTAL SETUP: OVERVIEW



EXPERIMENTAL SETUP: OVERVIEW



EXPERIMENTAL SETUP: OVERVIEW



Mean Absolute Percentage Error

EXPERIMENTAL SETUP: OVERVIEW

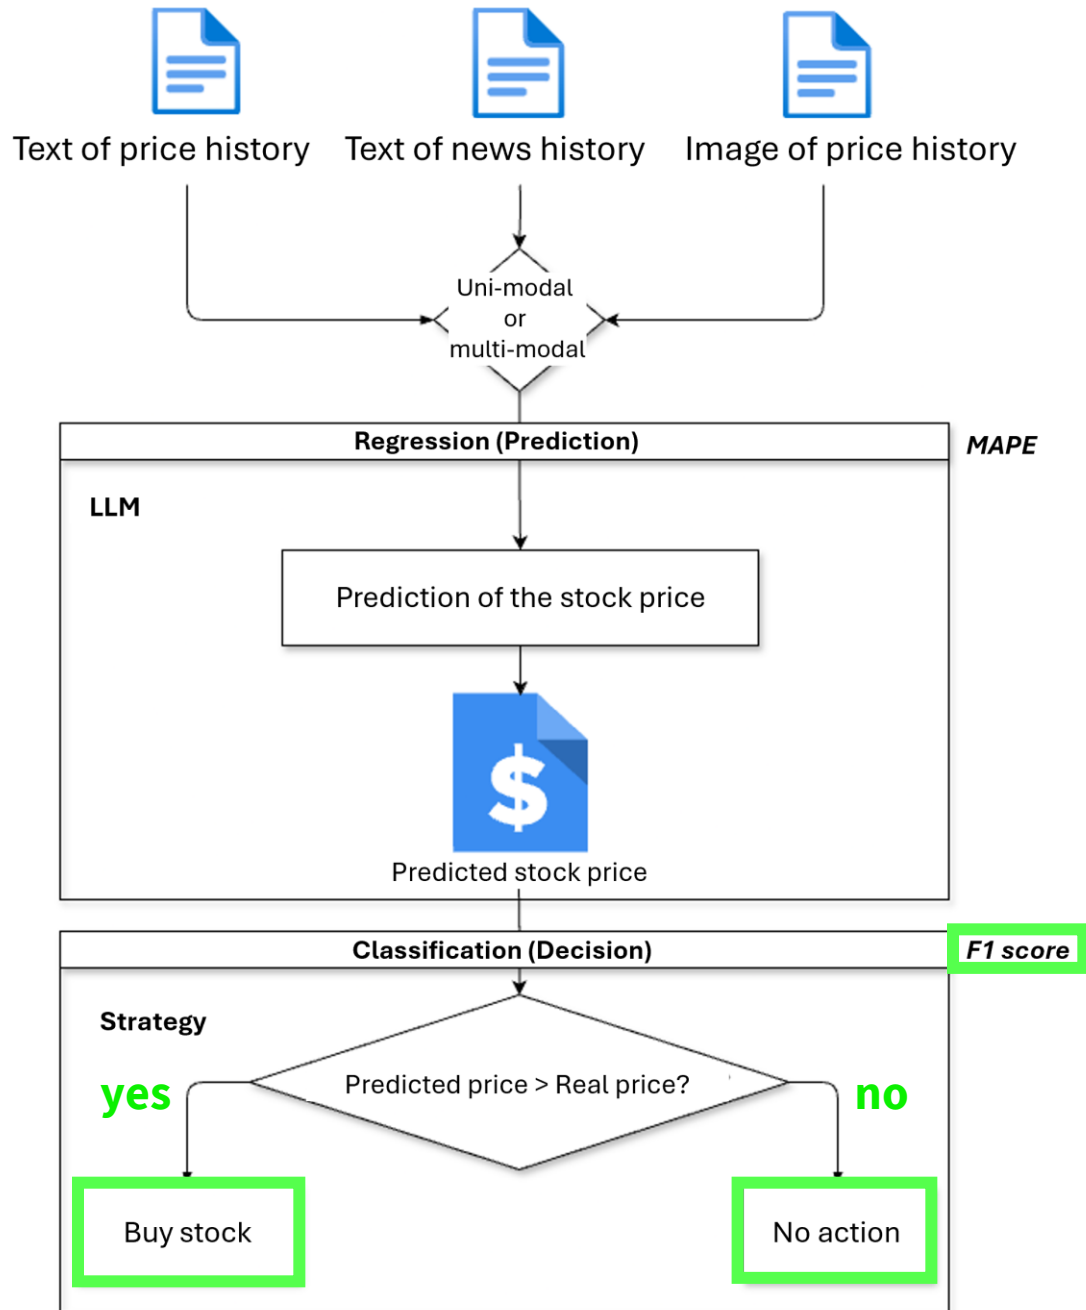


Image Source: Horn & Schlippe (2025).



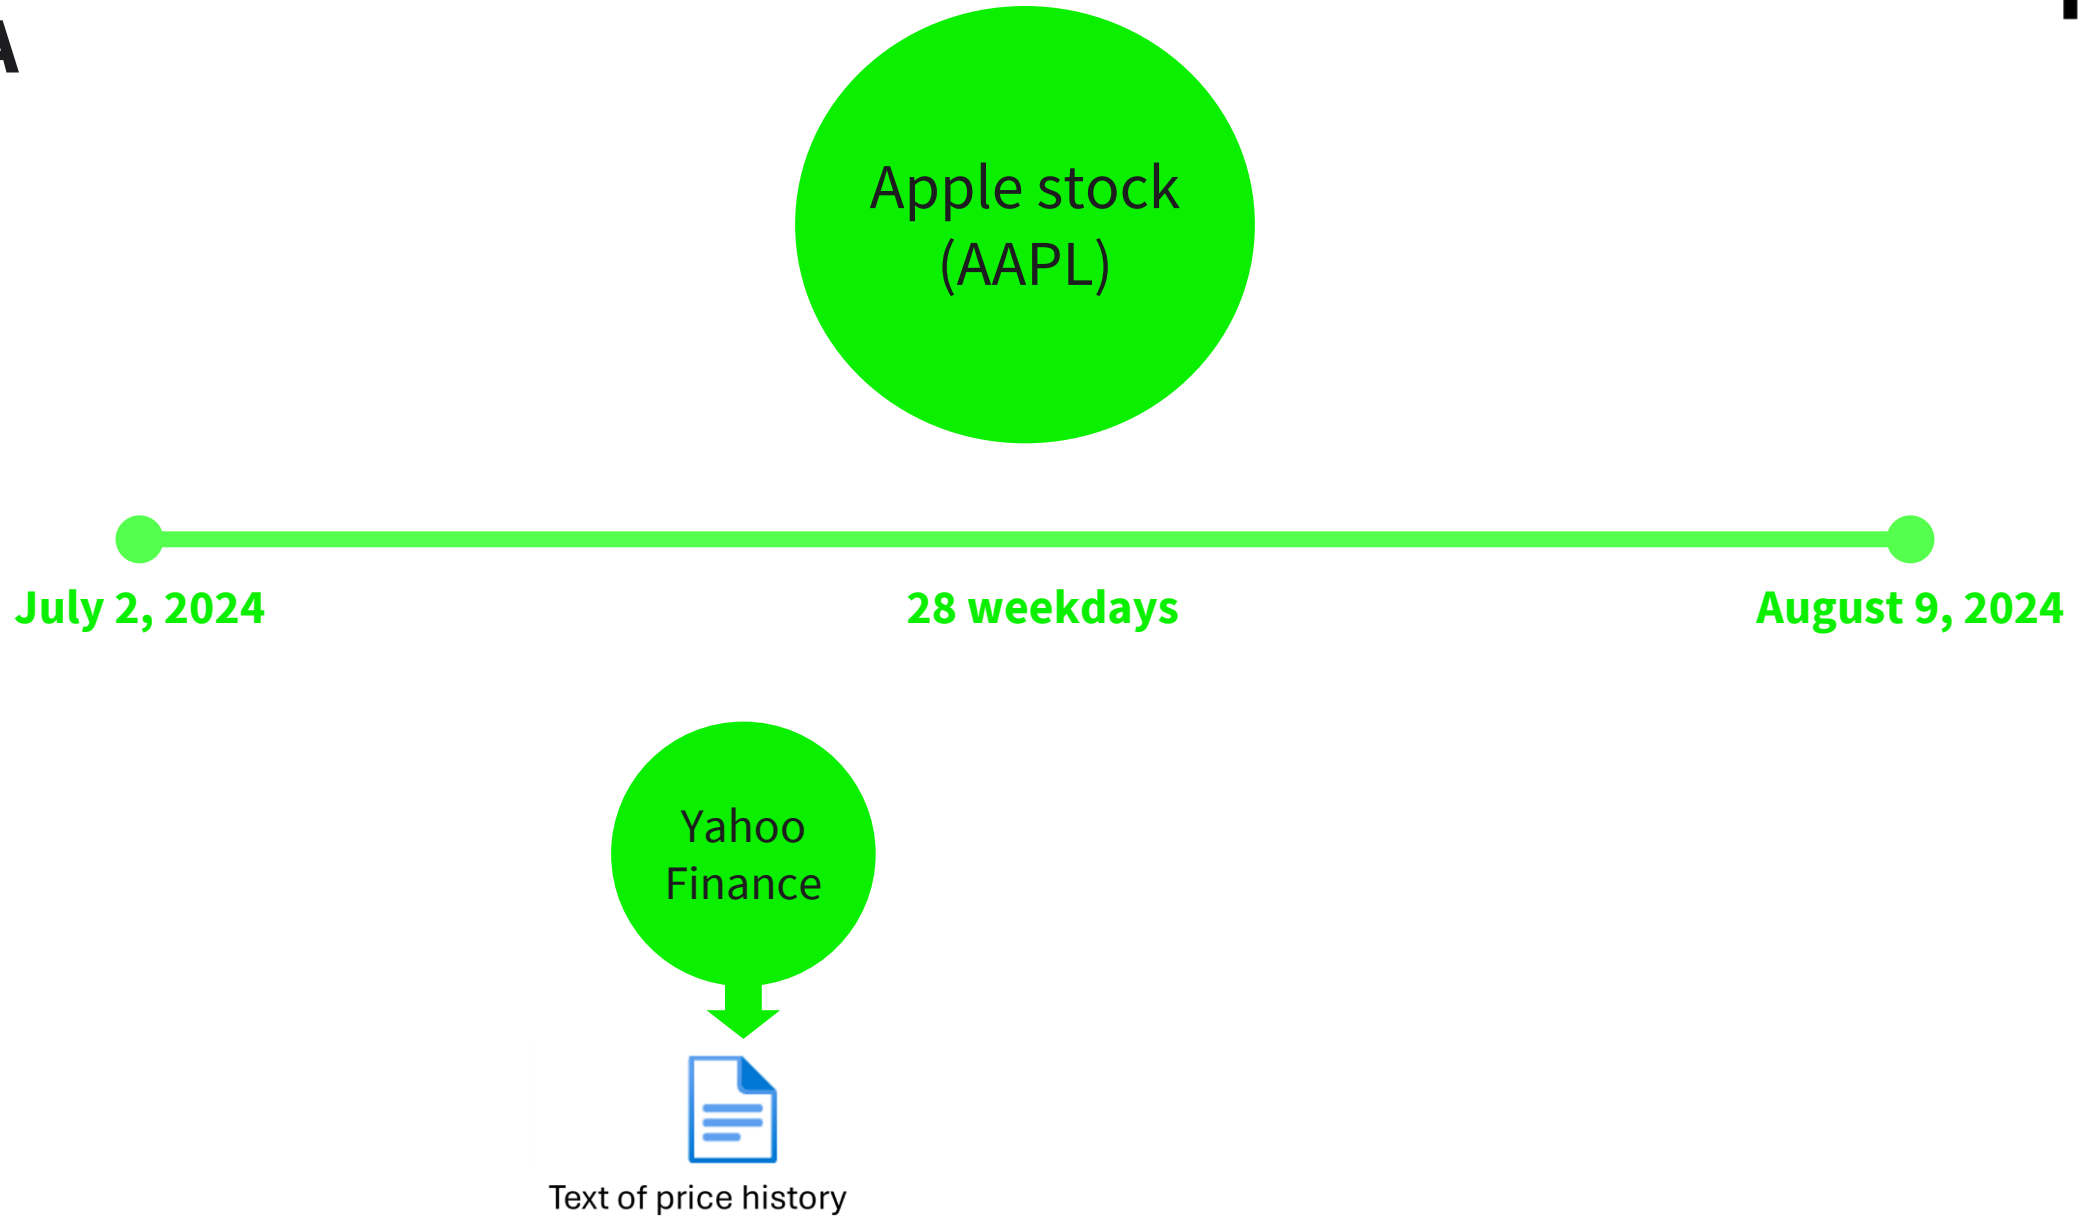
Text of price history

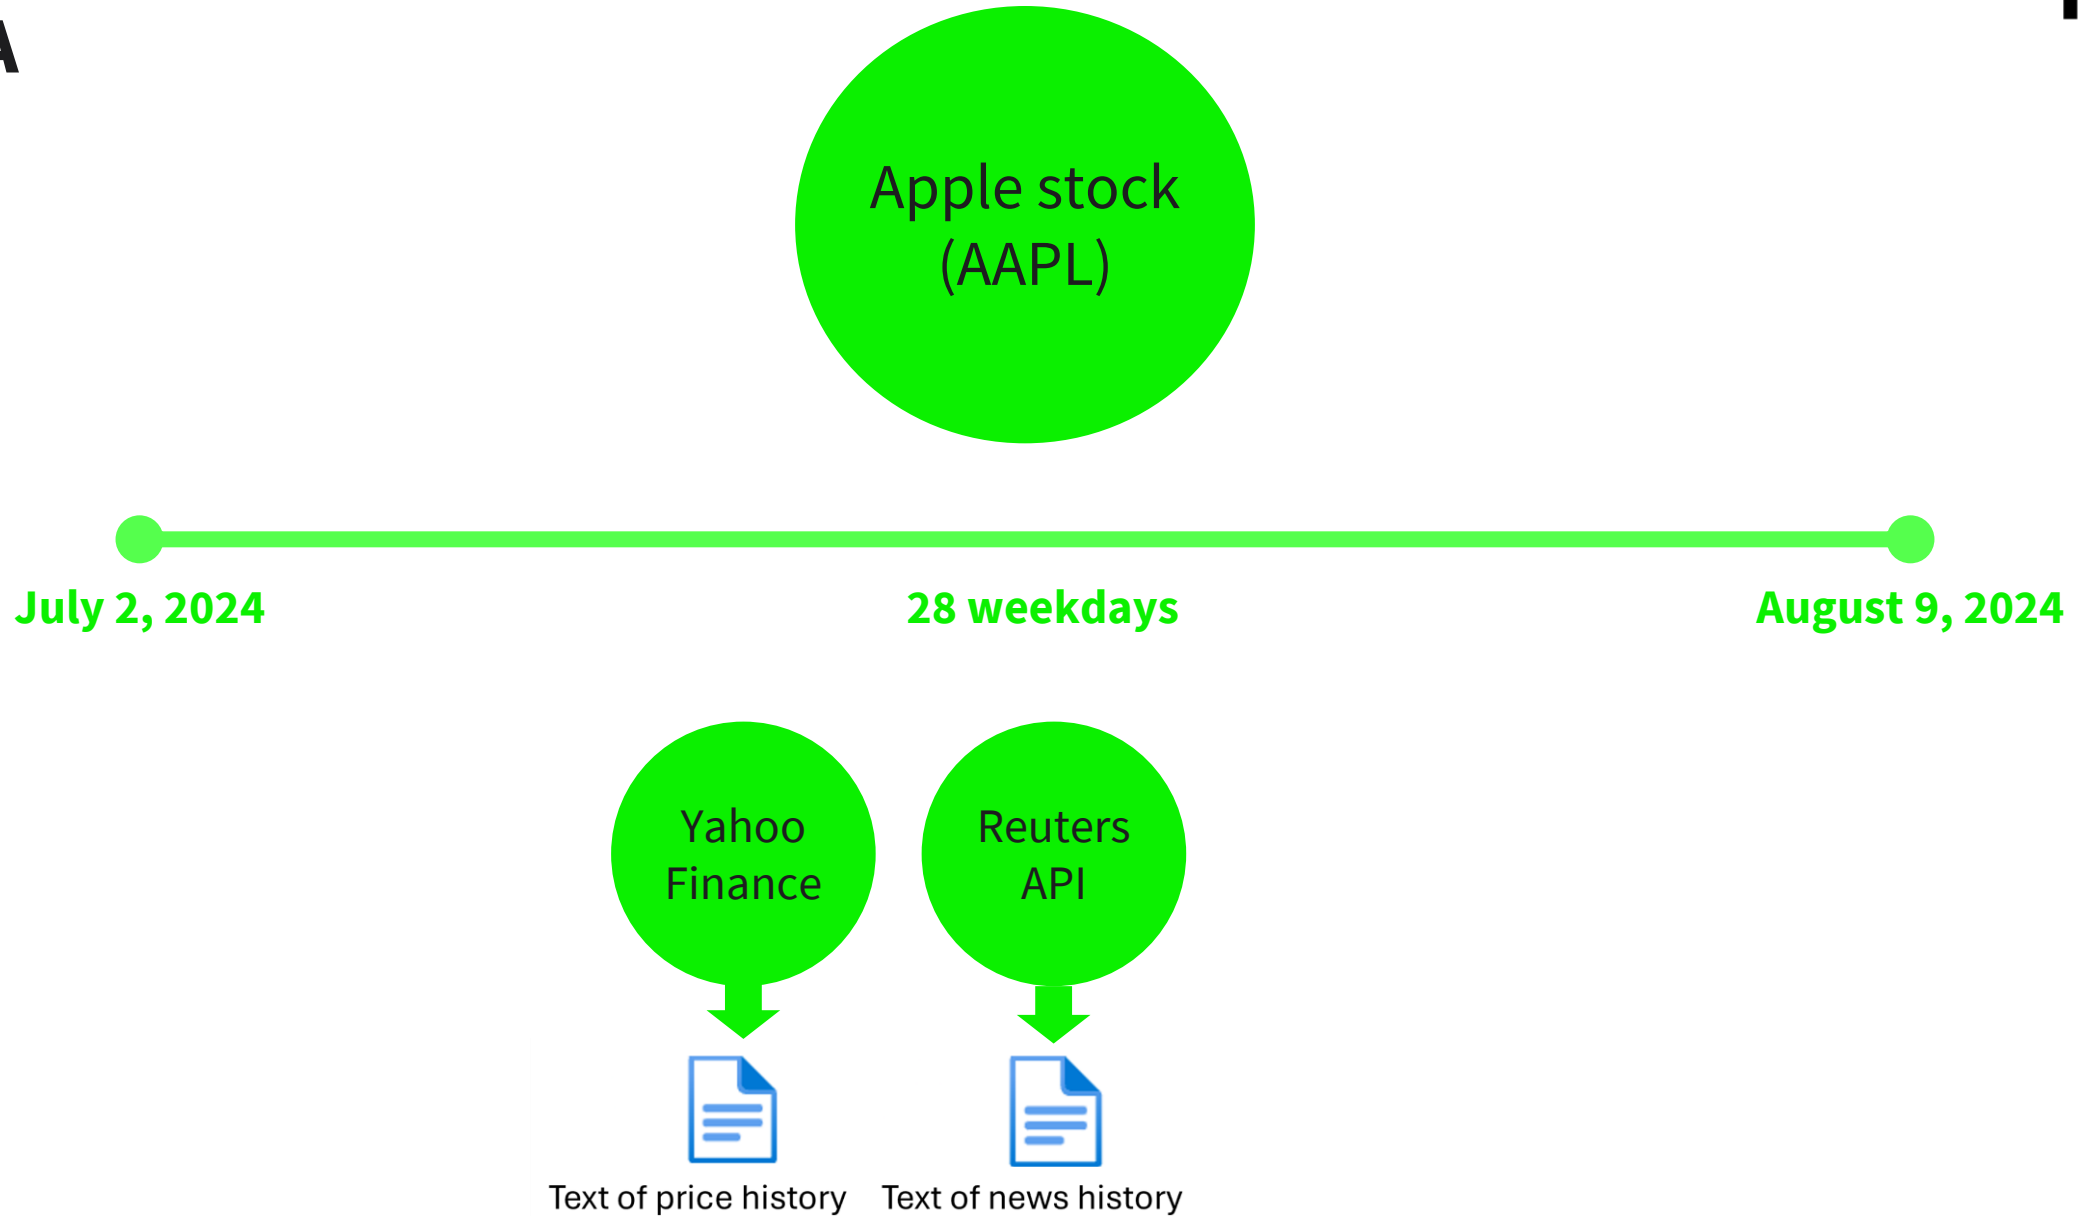


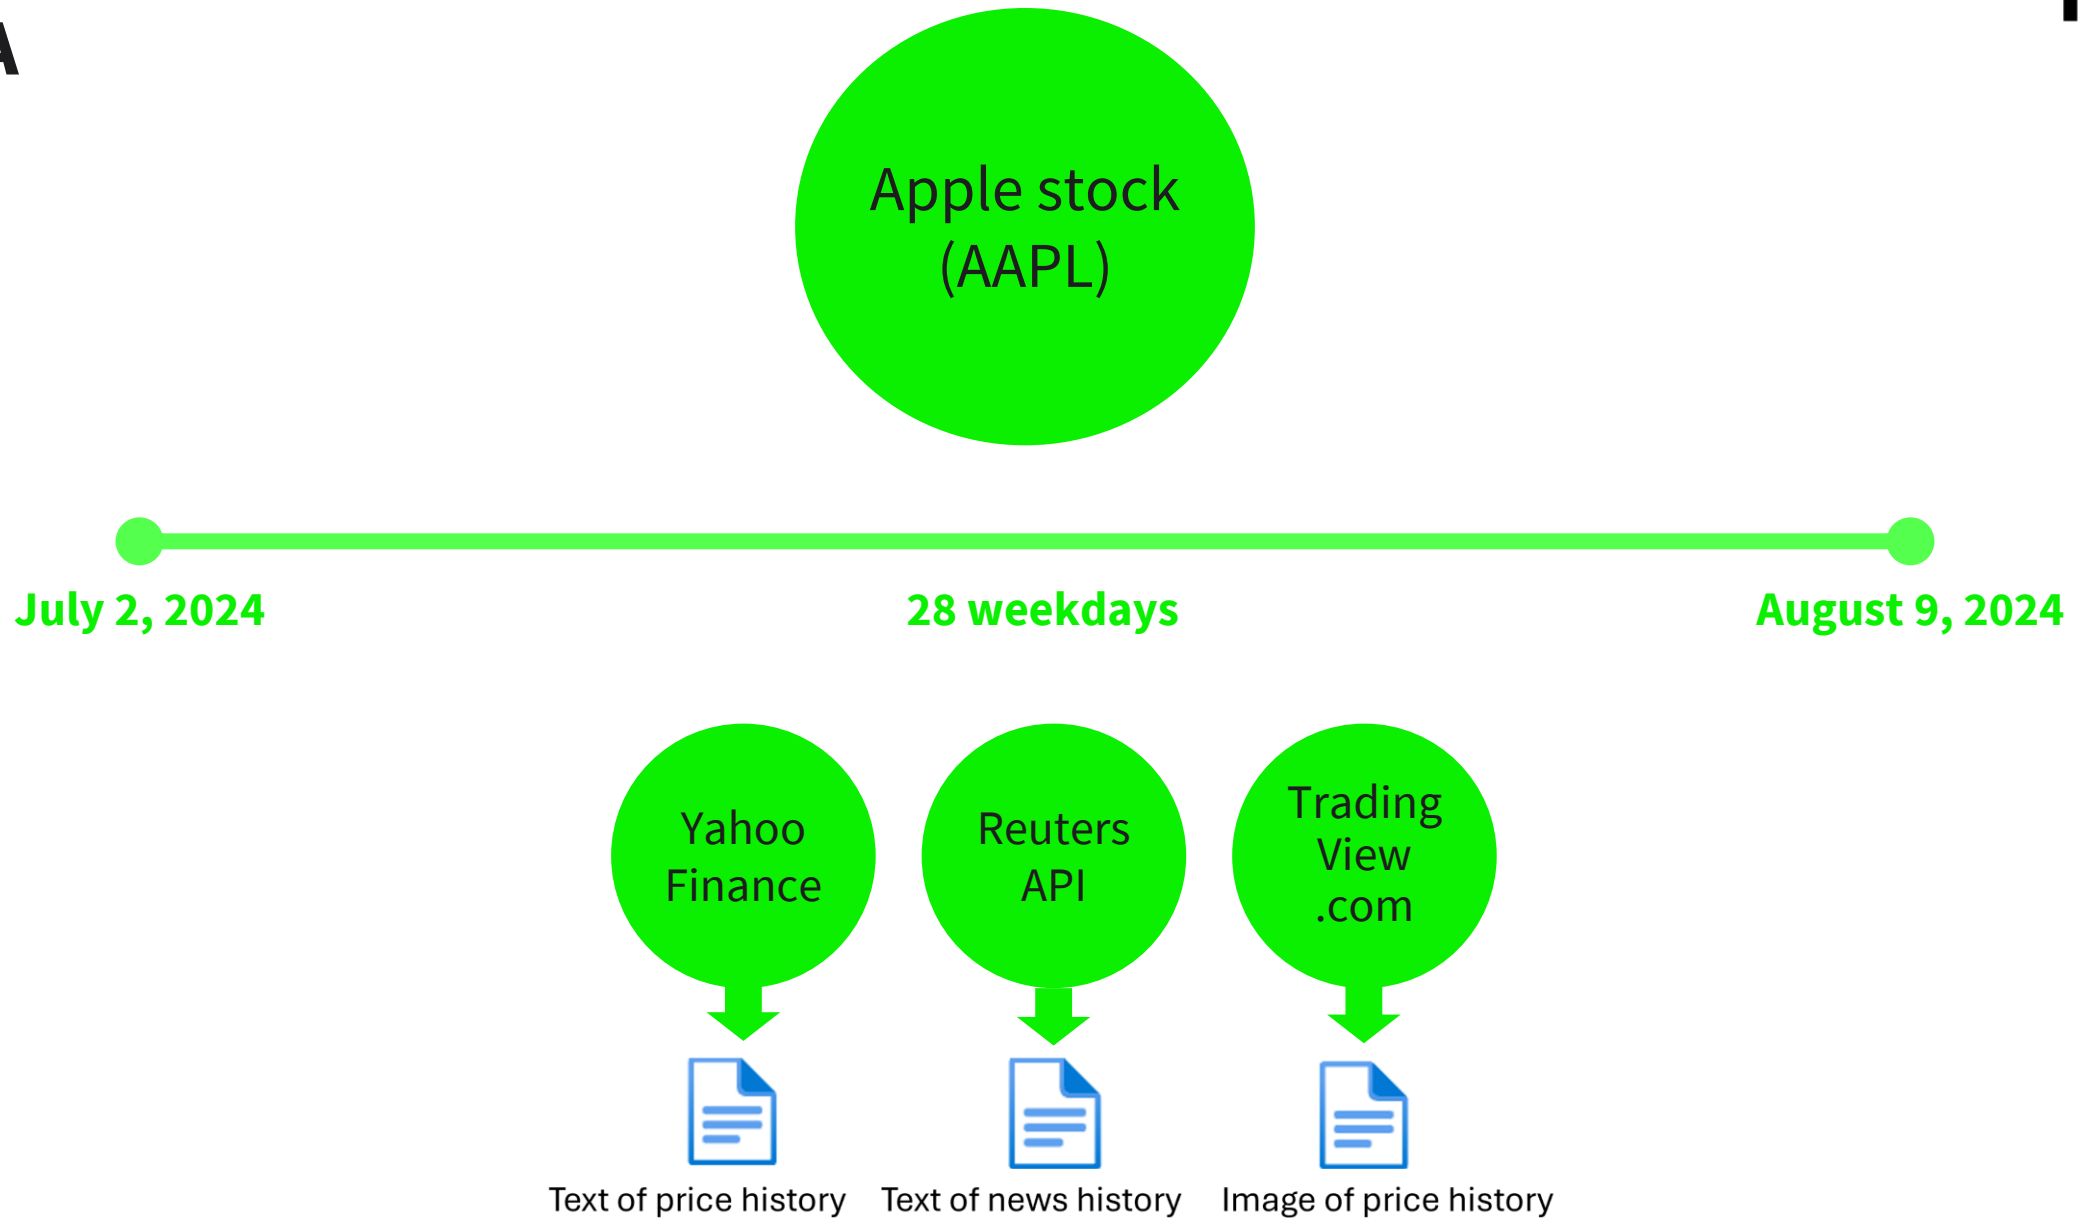
Text of news history



Image of price history







EXPERIMENTS AND RESULTS

PROMPTS

Prompt to Predict the Closing Price Leveraging a Text of the Price History:

Attached you will find a CSV file of the daily closing prices of the <SHARE NAME> share from the period of 1 year ago until today. The last closing price of the share was <PRICE>. Based on the information in the attached file, make a prediction of what the closing price will be today which is <DATE>. Please output only the predicted price in \$. Please refrain from any other outputs.

Prompt to Predict the Closing Price Leveraging a Text of the News History:

Attached you will find a file containing the last 20 news about the <SHARE NAME> share from the period of 1 year ago until today. Every news text contains the corresponding date, news title and the news content. The last closing price of the share was <PRICE>. Based on the information in the attached file, make a prediction of what the closing price will be today which is <DATE>. Please output only the predicted price in \$. Please refrain from any other outputs.

Prompt to Predict the Closing Price Leveraging an Image of the Price History:

Attached you will find an image file of the daily closing prices of the <SHARE NAME> share from the period of 1 year ago until today. The last closing price of the share was <PRICE>. Based on the information in the attached file, make a prediction of what the closing price will be today which is <DATE>. Please output only the predicted price in \$. Please refrain from any other outputs.

PROMPTS

Prompt to Predict the Closing Price Leveraging a Text of the Price History:

Attached you will find a CSV file of the daily closing prices of the <SHARE NAME> share from the period of 1 year ago until today. The last closing price of the share was <PRICE>. Based on the information in the attached file, make a prediction of what the closing price will be today which is <DATE>. Please output only the predicted price in \$. Please refrain from any other outputs.

Prompt to Predict the Closing Price Leveraging a Text of the News History:

Attached you will find a file containing the last 20 news about the <SHARE NAME> share from the period of 1 year ago until today. Every news text contains the corresponding date, news title and the news content. The last closing price of the share was <PRICE>. Based on the information in the attached file, make a prediction of what the closing price will be today which is <DATE>. Please output only the predicted price in \$. Please refrain from any other outputs.

Prompt to Predict the Closing Price Leveraging an Image of the Price History:

Attached you will find an image file of the daily closing prices of the <SHARE NAME> share from the period of 1 year ago until today. The last closing price of the share was <PRICE>. Based on the information in the attached file, make a prediction of what the closing price will be today which is <DATE>. Please output only the predicted price in \$. Please refrain from any other outputs.

PROMPTS

Prompt to Predict the Closing Price Leveraging a Text of the Price History:

Attached you will find a CSV file of the daily closing prices of the <SHARE NAME> share from the period of 1 year ago until today. The last closing price of the share was <PRICE>. Based on the information in the attached file, make a prediction of what the closing price will be today which is <DATE>. Please output only the predicted price in \$. Please refrain from any other outputs.

Prompt to Predict the Closing Price Leveraging a Text of the News History:

Attached you will find a file containing the last 20 news about the <SHARE NAME> share from the period of 1 year ago until today. Every news text contains the corresponding date, news title and the news content. The last closing price of the share was <PRICE>. Based on the information in the attached file, make a prediction of what the closing price will be today which is <DATE>. Please output only the predicted price in \$. Please refrain from any other outputs.

Prompt to Predict the Closing Price Leveraging an Image of the Price History:

Attached you will find an image file of the daily closing prices of the <SHARE NAME> share from the period of 1 year ago until today. The last closing price of the share was <PRICE>. Based on the information in the attached file, make a prediction of what the closing price will be today which is <DATE>. Please output only the predicted price in \$. Please refrain from any other outputs.

PROMPTS

Prompt to Predict the Closing Price Leveraging a Text of the Price History:

Attached you will find a CSV file of the daily closing prices of the <SHARE NAME> share from the period of 1 year ago until today. The last closing price of the share was <PRICE>. Based on the information in the attached file, make a prediction of what the closing price will be today which is <DATE>. Please output only the predicted price in \$. Please refrain from any other outputs.

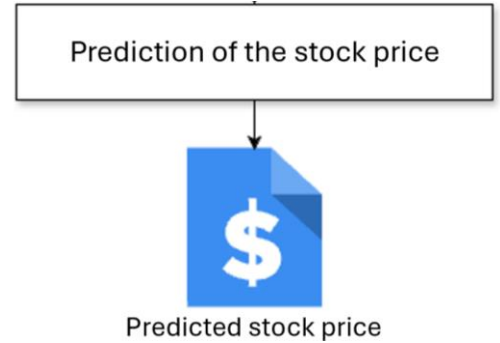
Prompt to Predict the Closing Price Leveraging a Text of the News History:

Attached you will find a file containing the last 20 news about the <SHARE NAME> share from the period of 1 year ago until today. Every news text contains the corresponding date, news title and the news content. The last closing price of the share was <PRICE>. Based on the information in the attached file, make a prediction of what the closing price will be today which is <DATE>. Please output only the predicted price in \$. Please refrain from any other outputs.

Prompt to Predict the Closing Price Leveraging an Image of the Price History:

Attached you will find an image file of the daily closing prices of the <SHARE NAME> share from the period of 1 year ago until today. The last closing price of the share was <PRICE>. Based on the information in the attached file, make a prediction of what the closing price will be today which is <DATE>. Please output only the predicted price in \$. Please refrain from any other outputs.

RESULTS: MEAN ABSOLUTE PERCENTAGE ERROR



MAPE



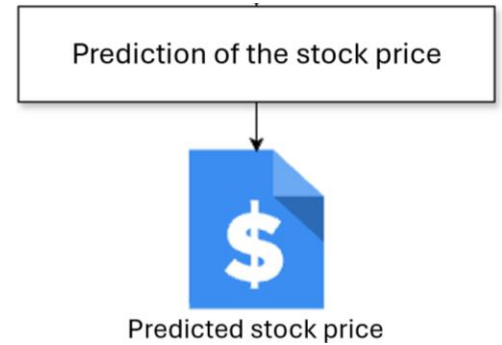


Image Source: Horn & Schlippe (2025).



RESULTS: MEAN ABSOLUTE PERCENTAGE ERROR




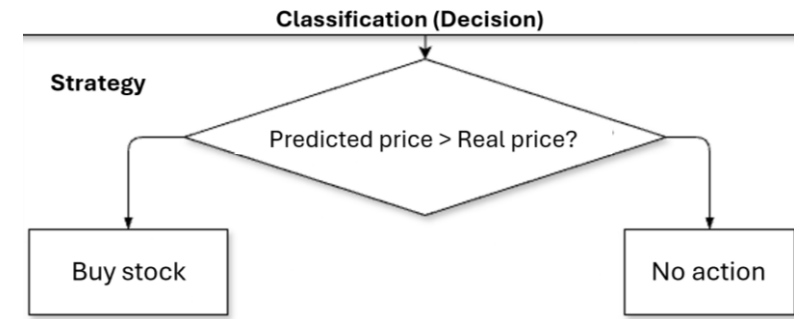
LOWEST MAPE: 1.4%

Image Source: Horn & Schlippe (2025).

RESULTS: F1 SCORE



RESULTS: F1 SCORE



HIGHEST F1: 83%

5

CONCLUSION AND FUTURE WORK

CONCLUSION AND FUTURE WORK

Conclusion

- LLMs leverage multimodal data for market movement predictions.

CONCLUSION AND FUTURE WORK

Conclusion

- LLMs leverage multimodal data for market movement predictions.
- Claude 3.5 and Gemma 2 achieved the lowest MAPE of 1.4% using only price history text.
 - Demonstrates the potential of simple input formats for accurate predictions.

CONCLUSION AND FUTURE WORK

Conclusion

- LLMs leverage multimodal data for market movement predictions.
- Claude 3.5 and Gemma 2 achieved the lowest MAPE of 1.4% using only price history text.
 - Demonstrates the potential of simple input formats for accurate predictions.
- Combining price history with news history improved performance for Claude 3.5 and Mistral 0.3.
 - Integration of diverse data sources enhances forecasting capabilities.

CONCLUSION AND FUTURE WORK

Conclusion

- LLMs leverage multimodal data for market movement predictions.
- Claude 3.5 and Gemma 2 achieved the lowest MAPE of 1.4% using only price history text.
 - Demonstrates the potential of simple input formats for accurate predictions.
- Combining price history with news history improved performance for Claude 3.5 and Mistral 0.3.
 - Integration of diverse data sources enhances forecasting capabilities.

Future Work

- Investigate other stocks
- Explore additional modalities
- Analyze LLM-based multi-agent scenarios

THANK YOU

Tim Schlippe

 tim.schlippe@iu.org