Using Large Language Models to Predict Crude Oil Price based on Financial Sentiment Analysis

Hind Aldabagh[1], Xianrong Zheng[2], Mohammad Najand[3], Sandeep Kalari[1], Ravi Mukkamala[1]
[1]Department of Computer Science
[2]Department of Information Technology & Decision Sciences
[3]Department of Finance
Old Dominion University, Norfolk, Virginia, USA
{halda001, x1zheng, mnajand, skala003, rmukkama}@odu.edu

***Abstract*—Financial sentiment analysis is essential for interpreting market dynamics and making informed strategic trading decisions. Although recent years have seen the adoption of advanced deep learning architectures and financial language models, this study advances the field by exploring the capability of large language model, specifically Gemini 2.5 pro offered by Google, for financial sentiment prediction, with a particular focus on the crude oil market. By monitoring news headlines and news content for the year 2025 and using a few-shot prompt technique, we evaluate multiple Gemini 2.5 prompt formulations on a carefully curated dataset of crude oil news headlines and text. The resulting prediction model is evaluated using the mean absolute error and the mean square error of the sentiment labels and the prediction prices. Incorporating news based textual information significantly improved forecasting performance. The RMSE decreased from 1.57 when using only historical price data to 1.14 when combining price information with news headlines and full article text through few-shot prompting. Similarly, the MAE improved from 1.32 to 0.78 under the same enhanced input configuration, demonstrating that incorporating news information substantially reduces the prediction error. These results suggest that the sentiments derived from the news provide valuable incremental information for short-term crude oil price prediction.**

***Index Terms*—AI, Financial Sentiment Analysis, Large Language Models, Crude Oil Price Prediction**

I. Introduction

Crude oil represents a crucial commodity in the global economy, with price changes that have far-reaching consequences for nations, businesses, organizations, and individuals around the world [1]. Consequently, achieving reliable forecasts of crude oil prices is essential for effective economic planning and informed decision-making processes at various levels.

The news headlines and associated textual information play a critical role in predicting crude oil prices, as they provide real-time information on geopolitical events, supply-demand dynamics, market speculation, and policy changes that can rapidly influence price movements [2]. Headlines can capture breaking developments such as sanctions, inventory reports, production decisions, or disruptions in key regions—factors that might not be immediately reflected in quantitative data, but can cause significant market volatility. By analyzing the language and sentiment of these news sources, predictive models can better anticipate changes in market expectations and respond more precisely to emerging risks and opportunities, ultimately improving the timeliness and precision of oil price forecasts [3].

Prompt engineering has recently emerged as a discipline dedicated to designing and refining input prompts to maximize the effectiveness of language models (LMs) across a wide range of analytical and decision-making tasks. The mastery of prompt engineering facilitates a deeper understanding of the strengths and limitations of large language models (LLMs), particularly when applied to data-driven financial problems such as forecasting crude oil prices [4].

In the context of financial research, prompt engineering is increasingly used to enhance the ability of LLMs' to process and interpret market-relevant information, including news events, economic indicators, geopolitical updates, and supply-and-demand dynamics. By carefully constructing prompts, researchers can guide LLMs toward more accurate sentiment extraction, improved reasoning over market signals, and more consistent generation of predictive insights [5].

Beyond improving raw output quality, prompt engineering serves as a critical framework to ensure that LLM-based forecasting systems behave reliably and transparently [6]. It enables the development of prompting strategies that incorporate domain-specific knowledge, integrate external tools (e.g., price databases or econometric models), or reinforce safety mechanisms that prevent hallucinations. In the domain of crude oil price prediction, prompt engineering is therefore not only a technique to improve the accuracy of the model, but a key enabler of robust and interpretable AI-assisted financial forecasting [7].

In this work, we explore large language model (LLM) prompt engineering techniques to predict the directional movement of WTI crude oil prices. Specifically, we investigate the impact of incorporating sentiment information extracted from news articles that are likely to influence oil prices. To this end, LLMs are employed in a two-step approach: first, to derive sentiment signals from the collected news items, and second, to predict future oil prices based on these sentiments. The methodology leverages a few-shot prompting strategy, in which the LLM is provided with labeled examples from news data spanning June to November 2025 through a customized prompt. The extracted sentiment, together with prior oil price data, is then used to generate next-day price predictions. The results demonstrate that including news-driven sentiment improves forecast accuracy compared to models that rely solely on historical price data. These findings suggest that LLMs can effectively extract valuable market-relevant information from textual data, improving predictive modeling of crude oil price movements. Improvement is measured in test data using RMSE and MAE as metrics.

The remainder of the paper is organized as follows. In Section 2, we briefly summarize previous work in this area. Section 3 describes the experimental methodology adopted, including data collection, the prompts used, and the test data.In Section 4, we discuss the results of the experiments. Finally, Section 5 summarizes the contributions of our work and our future plans.

II. Related Work

In this section, we briefly summarize recent developments in crude oil price prediction using traditional machine learning models, the use of sentiment analysis using LLMs, and prompt engineering.

A. Oil Price Prediction Using News and Traditional Machine Learning Methods

Traditional econometric and time-series models have been widely used to forecast crude-oil prices; however, recent studies show that integrating textual information from financial news substantially improves predictive accuracy [8]. With the increasing availability of real-time news streams, research has focused on the contribution of headline sentiment, topic structures, and event representations to short-term price movements [9].

Early work examined sentiment extracted from oil-related news. For example, [10] reported that sentiment indices derived from energy-market news improve forecasts of returns and volatility when combined with signal decomposition and deep neural networks. To address the sparsity of short headlines, [11] developed a domain-specific sentiment lexicon tailored to crude-oil markets, producing stronger predictions of price levels and directional shifts.

Another research direction integrates text-derived features with traditional numerical predictors. The authors in [2] demonstrated that combining textual features with historical prices in a boosting-based framework significantly improves the accuracy of short-term forecasting. Event-level representations have also proven effective; [12] showed that the features extracted from geopolitical disruptions, supply shocks, and policy announcements outperform generic sentiment indicators.

Recent work emphasizes domain-specific transformer models. CrudeBERT [13], trained in oil-market text, produces semantic representations more aligned with commodity-market narratives than general-purpose models. Topic-conditioned sentiment analysis has also shown promise; [14] found that topic-specific sentiment signals outperform aggregate sentiment measures.

Multimodal architectures now combine market data, sentiment indicators, event representations, and transformer embeddings. For example, [15] reported that multimodal deep-learning models outperform single-source approaches.

In general, the literature consistently concludes that the integration of textual information with historical price data improves the forecasting of crude-oil prices. Event extraction, topic-aware sentiment modeling, and transformer-based representations represent the most promising current directions.

B. Financial Sentiment Analysis Using Large Language Models

Recent studies increasingly employ large language models (LLMs) to extract such signals. The framework proposed in [16] uses pretrained transformer-based LLMs and demonstrates notable improvements over econometric and shallow machine-learning methods when modeling the highly nonlinear and volatile nature of crude-oil markets. Similarly, [17] integrates the sentiment and textual features derived from LLM with traditional time-series variables, achieving greater precision of the forecast.

In general, combining sentiment and contextual features obtained from LLMs with numerical predictors has emerged as an effective strategy to develop more adaptive and robust crude-oil forecasting models.

C. Prompt Engineering

Recent studies show that prompt design has a substantial impact on the performance of LLMs in financial applications. Zero-shot and few-shot prompting are the two dominant paradigms: prior work uses zero-shot instructions for headline classification, volatility tagging, and event detection, but reports highly variable accuracy in specialized domains such as energy and commodities [18]. Few-shot prompting, where a small number of labeled exemplars are embedded in the prompt, has been found to improve label consistency, directional accuracy, and calibration for tasks such as earnings-call sentiment, credit-risk assessment, and trading-signal extraction [19].

Beyond simple label templates, several authors propose structured prompts that decompose the task. For example, LLMs are first instructed to produce factor-based summaries of financial text (e.g., separating micro news, macro indicators, and policy shocks) and are then asked, in a second step, to infer sentiment or price direction from those structured summaries [20]. Other works incorporate brief rationales and confidence scores directly in the prompt to stabilize outputs across heterogeneous news categories (geopolitics, regulatory changes, and inventory reports) and to facilitate downstream use in forecasting pipelines [21].

Our approach follows this line of work by employing a few-shot, two-stage prompt tailored to oil-market news, first summarizing daily headlines and articles into market-relevant factors, and then eliciting a directional and numerical prediction of next-day WTI prices. We also compared the results of few-shot prompting with zero-shot prompting.

III. Experimental Methodology

A. Framework overview

The overall architecture of the proposed crude-oil price forecasting system is shown in Fig. 1. The framework includes two experimental settings. Setting 1 predicts the next-day crude-oil price using only historical price data via few-shot prompting. Setting 2 incorporates both price data and news text; the LLM first produces a summary of daily news and then predicts the next-day price using few-shot prompting. Performance across the two settings is evaluated using RMSE and MAE.

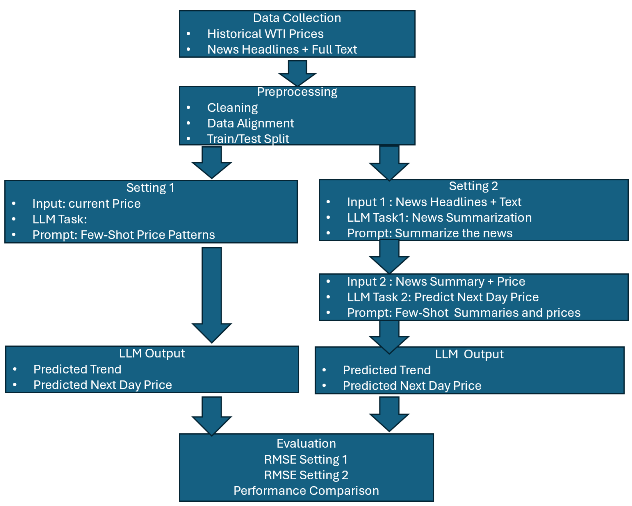

Fig. 1. Data flow for the two experimental settings: (1) price-only forecasting and (2) price-plus-news forecasting using few-shot prompting with a large language model.

B. Data Collection and Description

The news data was collected from OilPrice.com, a widely used source of real-time energy-market reporting. Each record includes a headline, full text of the article, publication timestamp, URL, and source. To capture both high-level cues and detailed context, the dataset combines: **Headlines**, which provide the primary narrative or sentiment, and **Full-length** articles, which supply richer details on market drivers and sentiment.

Data were gathered over a six-month period (June–November 2025), covering varying market conditions, including periods of significant volatility. Fig. 2 displays the WTI price trajectory during this interval.

The preprocessing steps included the removal of duplicates, normalization of encoding and punctuation, elimination of non-editorial content, and verification of chronological consistency. The final data set consists of approximately 70 pairs of headlines and articles. Ten representative few-shot examples used in the prompt are shown in Table I; they include both

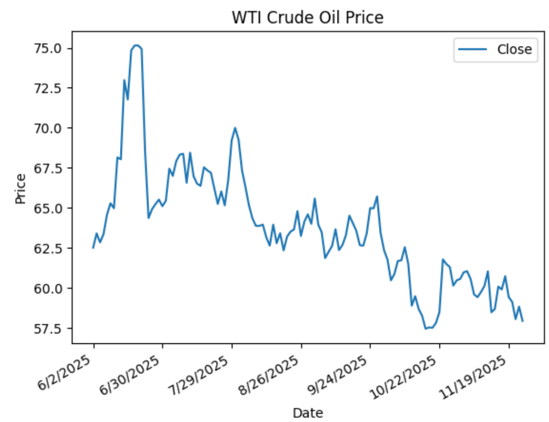

Fig. 2. Oil Prices.

upward and downward market signals selected to minimize ambiguity.

TABLE I
Oil Market News Headlines and Corresponding Dates

Date	Headline
2025-11-12	Crude Oil, Distillates Inventories See Minor Build as Prices Sink.
2025-11-12	Trump Signs Deal to Reopen U.S. Government
2025-11-12	U.S. Sanctions Widen Russia's Crude Discount to $20 a Barrels
2025-11-12	India's Top Refiner Looks to Buy Non-Sanctioned Russian Crude
⋮	⋮
2025-10-01	Iran-Aligned Houthis Sanction U.S. Oil Majors
2025-10-01	Crude Oil Plummets to Lowest Since June
2025-10-01	Global Investors Flock to Middle East Infrastructure Opportunities
⋮	⋮
2025-06-04	Oil Prices Extend Gains on Supply Worry
2025-06-04	Large Build in Gasoline Stocks Overshadows Losses in Crude Inventories
2025-06-04	The OPEC+ Nations That Were Reluctant to Boost Production

C. Prompt Engineering and Few Shots Prompting

Prompt engineering was applied systematically to optimize LLM behavior. Initial prompts were iteratively refined based on output accuracy, relevance, and consistency with economic logic. Both zero-shot and few-shot formats were considered, with few-shot prompting selected for the forecasting task due to its superior alignment with domain-specific semantics.

The composite prompt used for summarization and prediction contains two sequential roles:

Financial News Analyst, responsible for generating a structured summary of supply, demand, geopolitical, inventory, sentiment, and analyst-driven market factors; and

Price Prediction Analyst, responsible for forecasting direction, next-day price, reasoning, and confidence based on the generated summary and current price.

The full prompt is shown below for reproducibility.

SYSTEM INSTRUCTION: SUMMARIZATION OF CRUDE OIL NEWS AND PRICE PREDICTION

Role 1: Financial News Analyst (Crude Oil). You are a financial news analyst specialized in crude oil markets.

Inputs. You receive: (1) a calendar date and (2) a set of crude-oil-related news articles, each with headline, source, and full text.

Task. Read all articles and produce a concise and structured summary of information that is relevant to predict the *next-day* price of crude oil. Focus only on: (i) supply factors (OPEC/OPEC+ decisions, outages, production changes, sanctions, export restrictions),
(ii) demand factors (global growth, macroeconomic data, travel/-mobility, industrial activity),
(iii) geopolitical risks (wars, conflicts, shipping disruptions, sanctions or sanctions threats),
(iv) inventories and storage (EIA/API reports, stock builds/draws, floating storage, strategic reserves),
(v) market sentiment (risk-on/risk-off tone, explicitly bullish or bearish commentary on oil), and
(vi) analyst views (explicit forecasts or price targets for crude oil).

Role 2: Price Prediction Analyst. You are a crude oil market analyst who learns from past patterns in news summaries and price movements. You are given several past examples; for each example, you see (a) the structured summary of the day's crude oil news, (b) the today's crude oil price, and (c) the next-day price outcome.

Current task. Given the current day's structured news summary and today's crude oil price, produce the following outputs:

1) **Direction:** predict whether the next-day price will move UP or DOWN relative to today's price.
2) **Price:** predict the numeric next-day price in USD.
3) **Reasoning:** provide a brief explanation that links the news summary and today's price to your prediction.
4) **Confidence:** provide a confidence score between 0 and 1, where 0.2 indicates very low confidence, 0.5 medium confidence, and 0.8 or higher high confidence.

D. Test Data

The test data set consists of dates that follow those in the few-shot examples. Table II lists the corresponding headlines. This period features heightened uncertainty driven by geopolitical tensions, expanded U.S. sanctions on Russian crude, stranded cargoes, and competition for energy assets. At the same time, strong import activity from China and India, along with large-scale energy-sector investments, signals resilient demand. Additional reports highlight concerns about oversupply, rising inventories, and warnings of a potential oil glut. In general, headlines depict a market influenced by opposing forces: geopolitical risks supporting prices and supply growth combined with sanctions-induced trade realignments exerting downward pressure. At the same time, China and India appear to be increasing imports and strengthening energy partnerships, contributing to resilient demand despite market instability. Several articles highlight investment activity and large-scale infrastructure projects in Africa, the Middle East, and North America, suggesting long-term sustained confidence in the energy sector. Meanwhile, oil prices fluctuate in reaction to political developments, including peace negotiations, drilling policy disputes, and efficiency targets. In addition, concerns about oversupply, rising inventories, and warnings of a potential oil glut underscore the risk of downward price pressure. Overall, the headlines suggest a market shaped by competing forces geopolitical risk supporting prices, while

TABLE II
TEST DATA: OIL MARKET NEWS HEADLINES AND CORRESPONDING DATES

Date	Headline
2025-11-24	Florida Pushes Back Against Offshore Leasing Plan
2025-11-24	Nigeria's National Oil Company Records 64% Jump In Profits To $3.6B
2025-11-24	Russian Tanker Reaches Venezuela After Evading U.S. Warship
2025-11-21	Trump Outrages California Authorities With Oil Drilling Plan
2025-11-21	New U.S. Sanctions Could Leave 48 Million Barrels of Russian Oil Adrift at Sea
2025-11-21	Abu Dhabi Conglomerate Enters Bidding War for Lukoil Assets
2025-11-20	Oil Prices Sink After Zelenskyy Agrees to Work on Peace Deal
2025-11-20	Rosneft's Vanishing Dividend: A Warning Shot for Moscow's Oil Economy
2025-11-20	India's Reliance Industries Officially Shuts the Door on Russian Crude
2025-11-20	Controversial $5 Billion EACOP Project Is Now Three-Quarters Complete
2025-11-20	Angola Looks to Accelerate Mining Approvals to Attract Investment
2025-11-20	China's Oil Imports Surge as Middle East Flows Hit New Highs
2025-11-20	IEA Warns the World Is Falling Behind on Critical Energy Efficiency Targets
2025-11-19	Oil Prices Edge Higher as Traders Brace for Russian Sanctions Deadline
2025-11-19	A Trend Break? EIA Sees Jump in Alaska Oil Production in 2026
2025-11-19	Pakistan to Launch Artificial Offshore Island To Jump-Start Oil Exploration
2025-11-18	China's Oil Stockpiling Accelerated in October
2025-11-18	Reliance Snaps Up Kuwaiti Crude as India Shuns Sanctioned Russian Oil
2025-11-18	Sanctions Slow Russian and Iranian Crude Flows to China
2025-11-18	Turkish State Oil Giant Launches Landmark $4 Billion Islamic Bond
2025-11-18	Goldman Sachs: Oil Prices To Drop to $53 In 2026
2025-11-18	US Crude Oil Inventories Continue to Rise as Production Hits New High
2025-11-18	Exxon Joins Chevron in Eyeing Lukoil's Global Fire Sale
2025-11-17	OPEC+ Poised to Keep Pumping Despite Rising Oversupply Fears
2025-11-17	China and India Boost Crude Purchases as Global Glut Looms
2025-11-17	Kazakhstan Drags Big Oil to Swiss Court in $166 Billion Lawsuit
2025-11-17	Poland's Oil Giant Orlen Plugs Into EV Charging Market
2025-11-14	Canada and India to Partner in Critical Minerals Supply Chain
2025-11-14	Ukraine Intensifies Campaign Targeting Russia's Vital Oil Lifeline
2025-11-14	Iran Seizes Oil Tanker in Gulf of Oman
2025-11-14	Enbridge Moves to Rewrite North America's Heavy-Crude Map
2025-11-13	Libya Ships First-Ever Cargo From Long-Stalled Chadar Oil Field
2025-11-13	Hungary's MOL Moves to Acquire Stake in Sanctioned Serbian Refinery
2025-11-13	Norway's Oil Industry Raises 2026 Investment Forecast
2025-11-13	IEA Warns Oil Glut Will Be Worse Than Expected
2025-11-13	U.S. Sanctions Strand a Third of Russia's Crude Exports at Sea
2025-11-13	Oil Stabilizes After Selloff Amid OPEC Reassessment and U.S. Funding Deal

TABLE III
RMSE AND MAE COMPARISON BETWEEN PRICE-ONLY AND PRICE+NEWS FORECASTING SETTINGS

Model Setting	Inputs Used	RMSE	MAE
Setting 1	Few-Shot Price only	1.57	1.32
Setting 2	Few-Shot Price + News	**1.14**	**0.78**
Setting 3	Zero-Shot Price only	1.34	1.10
Setting 4	Zero-Shot Price + News	1.79	1.54

supply growth and sanctions-driven trade redirection exert countervailing downward pressure.

IV. RESULTS AND DISCUSSION

To assess the effect of incorporating textual information in a few-shot prompting for crude-oil price forecasting, two experimental configurations were evaluated. In the first, the model predicted the next-day price using only historical price data. In the second, the model utilized both the current price and LLM-generated summaries of contemporaneous news headlines and article text. Forecast accuracy was measured using the root mean squared error (RMSE) and mean absolute error (MAE). The price-only model achieved an RMSE of 1.57 and MAE of 1.32, while the model augmented with news content produced a substantially lower RMSE of 1.14 and MAE of 0.78.

This reduction in error demonstrates that the incorporation of contextual information from the news materially improves the predictive performance. Crude-oil prices are heavily influenced by geopolitical shocks, supply disruptions, sanctions, OPEC decisions, and large-scale market actions—factors that cannot be captured by numerical price series alone. When supplied with relevant textual signals, the model more effectively anticipated short-term reactions to market-moving events. The improved RMSE and MAE thus confirm that news headlines and article summaries provide incremental predictive value beyond historical price movements.

To further assess the effectiveness of the few-shot prompting, the RMSE and MAE obtained under this configuration were compared with the corresponding results from the zero-shot prompting setup. In this comparison, Settings 1 and 2 represent the few-shot prompting conditions, whereas Settings 3 and 4 correspond to the zero-shot prompting conditions. This evaluation enables a direct quantification of the performance gains attributable to the inclusion of few-shot examples.

In general, these results underscore the benefit of integrating structured numerical inputs with unstructured textual information in LLM-based forecasting frameworks. By combining news-driven sentiment with price dynamics, the model yields more accurate and responsive next-day predictions that better reflect real-world market behavior. Table III reports the RMSE and MAE values for the four settings, and Table IV compares the predicted and actual prices for all test days of the few-shot prompting.

The results indicate that the model achieves its highest forecasting accuracy when provided with few-shot examples that incorporate both current day prices and news information. In contrast, the poorest performance is observed in the zero-shot setting when the model relies solely on the current market price and news. This finding suggests that guided prompting with few-shots is essential for effective learning in this context, as the model appears to be susceptible to misinterpretation when exposed only to isolated daily news. From a financial sentiment analysis perspective, single-day news does not provide sufficient contextual grounding, whereas few-shot examples help reduce ambiguity and anchor the model toward meaningful market patterns.

TABLE IV
MODEL PREDICTION RESULTS FOR WTI CRUDE OIL

Prediction Date	Actual Price	Predicted Price (Setting 1)	Predicted Price (Setting 2)
2025-11-14	60.09	57.54	**59.64**
2025-11-15	59.91	**61.24**	62.24
2025-11-18	60.74	**61.46**	58.71
2025-11-19	59.44	61.89	**59.04**
2025-11-20	59.14	60.15	**59.84**
2025-11-21	58.06	58.25	**57.99**
2025-11-24	58.84	59.16	**58.67**
2025-11-25	57.95	59.94	**58.06**

These findings suggest that the integration of structured numerical price data with unstructured textual information improves predictive accuracy and supports the conclusion that the news sentiment contains incremental and measurable predictive power in short-term crude oil price forecasting.

A. *Analysis and Output Generation*

During the examined period, the LLM incorporated both supply–demand fundamentals and geopolitical risks to predict the next-day crude-oil price movements. Its reasoning captured the interaction between structural bearish forces—such as anticipated oversupply and rising inventories—and short-term bullish shocks arising from geopolitical events. The evolution of dominant market narratives is summarized below.

On 2025-11-14, the previous bearish sentiment from OPEC and IEA surplus forecasts had largely been priced in. However, a sudden disruption of 1.4 million bpd of Russian crude due to sanctions introduced immediate supply tightness, triggering a short-coverage rebound despite longer-term bearish pressures.

On 2025-11-15, the bullish momentum intensified after an attack on Russia's Novorossiysk port and Iran's seizure of a tanker near the Strait of Hormuz. The combination of a real supply loss and an elevated geopolitical risk at a critical choke point reinforced the upward pressure on prices.

By 2025-11-18, the narrative shifted bearish. Consensus forecasts projected a growing global surplus driven by increases in OPEC+ and non-OPEC output and downward revisions in global demand. Strong purchases from China and India were insufficient to offset expected supply growth.

On 2025-11-19, bearish sentiment deepened as record-high U.S. production and a 4.4-million-barrel API inventory build signaled immediate oversupply. A high-profile bearish forecast from Goldman Sachs further amplified downward expectations, even as sanctions continued to disrupt crude flows.

On 2025-11-20, the market sentiment somewhat balanced out. Downward pressure from a surprise EIA inventory build and weak demand was partially offset by near-term geopolitical risk from an impending U.S. sanctions deadline on Russian crude. A modest rebound was anticipated as traders priced in short-term supply tightening.

On 2025-11-21, opposing narratives persisted. Supply disruptions due to sanctions supported prices, while optimism about a possible peace deal exerted bearish influence. The disruption of 48 million barrels established a price floor, leading to a moderate rebound focused on near-term tightness.

By 2025-11-24, news of potential Ukraine peace negotiations reduced geopolitical risk, causing the market to shed part of its risk premium. This bearish development outweighed supportive factors, including a surprise U.S. inventory draw and confirmed OPEC+ supply increases for December.

Finally, on 2025-11-25, the market weighed two conflicting factors: a supply surge from Nigeria and unresolved tensions in Venezuela. Historical trends suggest that confirmed supply increases exert stronger and faster downward pressure than uncertain geopolitical disruptions, resulting in a modest price decline despite partial support from geopolitical risks.

V. Conclusion

News headlines and other textual information play a critical role in the prediction of crude-oil prices because oil markets react quickly to geopolitical events, macroeconomic announcements, supply–demand disruptions, and policy interventions. Unlike historical price data, which reflect only past behavior, news sources provide forward-looking signals that capture market expectations, investor sentiment, and real-time uncertainty. Headlines often report influential events—such as production cuts, sanctions, conflicts, and OPEC decisions—before their effects appear in trading activity or price movements. Consequently, incorporating textual information enables forecasting models to capture market-moving factors that are otherwise invisible to numerical time-series input.

Advances in natural language processing and large language models further strengthen this capability by improving sentiment extraction, identifying risk-relevant cues, and quantifying narratives embedded in news streams. By integrating these qualitative insights with traditional quantitative indicators, predictive systems achieve more accurate, robust, and adaptive performance in an environment characterized by volatility and rapid information dissemination.

In this paper, we employ LLMs to capture sentiments from the news headlines and the detailed news text to improve future crude oil predictions. Adding sentiments improved the accuracy of price prediction. We have used few-shot examples to train the LLM. The preliminary results are very encouraging.

We propose to extend this work by using agentic AI with a consortium of different LLM to extract the sentiments from the news and integrate them to a consistent set of sentiments using a reasoning LLM. The output of this LLM can then be used to predict prices using a predictive LLM.

References

[1] A. Tramboo and S. ul Rehman, "Trends in crude oil price and its economic impact," *South Asian Journal of Marketing & Management Research*, vol. 11, no. 5, 2021.

[2] Y. Bai, X. Li, H. Yu, and S. Jia, "Crude oil price forecasting incorporating news text," *International Journal of Forecasting*, vol. 38, no. 1, pp. 367–383, 2022.

[3] R. Lakatos, G. Bogacsovics, and A. Hajdu, "Predicting the direction of the oil price trend using sentiment analysis," in *2022 IEEE 2nd Conference on Information Technology and Data Science (CITDS)*, 2022, pp. 177–182.

[4] T. Pham and C. D. Cao, "Co-step: A prompt engineering framework improving llm's response," in *2025 10th International Conference on Applying New Technology in Green Buildings (ATiGB)*, 2025, pp. 588–593.

[5] Q. Lei and T. Zhu, "Qf-llm: Financial sentiment analysis with quantized large language model," in *AIDF '25: Proceedings of the 2025 International Conference on Artificial Intelligence and Digital Finance*, 2025, pp. 15–19.

[6] B. Chen, Z. Zhang, N. Langrené, and S. Zhu, "Unleashing the potential of prompt engineering for large language models," *Patterns*, vol. 6, no. 6, 2025.

[7] H. Ylinen, "Forecasting short-term crude oil price using global news," in *Master of Science Thesis, Tampere University*, 2025, p. 67.

[8] T. Yao, P.-F. Zhang, and Y.-J. Zhang, "An overview of crude oil price forecasting based on big data technology," *Frontiers of Engineering Management*, vol. 12, no. 4, pp. 938–951, 2025.

[9] S. Beyer Díaz, K. Coussement, A. De Caigny, L. F. Pérez, and S. Creemers, "Do the us president's tweets better predict oil prices? an empirical examination using long short-term memory networks," *International Journal of Production Research*, vol. 62, no. 6, pp. 2158–2175, 2024.

[10] Y. Zhang, H. Sun, and J. Wang, "The role of news sentiment in oil futures returns and volatility forecasting: A data-decomposition-based deep learning approach," *Energy Economics*, vol. 97, 2021.

[11] J. Chen, R. Liu, and S. Zhou, "Crude oil price fluctuation forecasting incorporating news sentiment based on an improved sentiment lexicon," *Energy Systems*, 2025.

[12] F. Liu and W. Huang, "Forecasting crude oil price using event extraction," *Energy Economics*, 2021.

[13] A. Jain and V. Kumar, "Crudebert: Applying economic theory towards fine-tuning transformer-based sentiment models for the crude oil market," *arXiv preprint arXiv:2305.06140*, 2023.

[14] Z. Wei, L. Zhang, and M. Xie, "Crude oil price forecasting with online news topic distribution and topic-based sentiment," *China Journal of Economics*, 2023.

[15] E. Ahlgren and L. Peterson, "Multimodal fusion of market data and news text for crude oil price forecasting using deep learning," *Energy AI*, 2024.

[16] K. He, L. Yu, and Y. Zou, "Crude oil future price forecasting using pretrained transformer model," *Procedia Computer Science*, vol. 242, pp. 288–293, 2024, 11th International Conference on Information Technology and Quantitative Management (ITQM 2024).

[17] K. He, Y. Li, Z. Peng, and Y. Zou, "Prediction of crude oil price using llm: An empirical analysis," *Procedia Computer Science*, vol. 266, pp. 594–602, 2025.

[18] G. Yong, K. Jeon, D. Gil, and G. Lee, "Prompt engineering for zero-shot and few-shot defect detection and classification using a visual-language pretrained model," *Computer-Aided Civil and Infrastructure Engineering*, 2022.

[19] M. R. Sudarshan and C. Kumbhar, "Comparative analysis of zero-shot, few-shot, and one-shot learning in natural language processing," in *2024 International Conference on Augmented Reality, Intelligent Systems, and Industrial Automation (ARIIA)*, 2024, pp. 1–6.

[20] S. Parmar, Y. Ganar, L. Henriques, and P. Thakre, "Comparative analysis of zero-shot, few-shot, and oneshot learning in machine learning," *International Journal of Scientific Research & Engineering Trends*, vol. 11, no. 2, 2025.

[21] A. Kim, M. Muhn, and V. Nikolaev, "Financial statement analysis with large language models," *arXiv preprint arXiv:2407.17866*, 2024.