

# ElliottAgents: A Natural Language-Driven Multi-Agent System for Stock Market Analysis and Prediction

**Jarosław A. Chudziak**

Warsaw University of Technology  
Institute of Computer Science  
Warsaw, Poland  
0000-0003-4534-8652  
jaroslaw.chudziak@pw.edu.pl

**Michał Wawer**

Warsaw University of Technology  
Institute of Control and Computation Engineering  
Warsaw, Poland  
0009-0004-2717-1616  
michal.wawer.stud@pw.edu.pl

## Abstract

This paper presents ElliottAgents, a multi-agent system leveraging natural language processing (NLP) and large language models (LLMs) to analyze complex stock market data. The system combines AI-driven analysis with the Elliott Wave Principle to generate human-comprehensible predictions and explanations. A key feature is the natural language dialogue between agents, enabling collaborative analysis refinement. The LLM-enhanced architecture facilitates advanced language understanding, reasoning, and autonomous decision-making. Experiments demonstrate the system's effectiveness in pattern recognition and generating natural language descriptions of market trends. ElliottAgents contributes to NLP applications in specialized domains, showcasing how AI-driven dialogue systems can enhance collaborative analysis in data-intensive fields. This research bridges the gap between complex financial data and human understanding, addressing the need for interpretable and adaptive prediction systems in finance.

## 1 Introduction

The integration of LLMs into multi-agent systems has opened new frontiers in AI, particularly in the domain of financial analysis (Zhao et al., 2023; Weng, 2024). Stock market prediction, a field characterized by its complexity and dynamism, has long challenged traditional AI-based methods. These approaches often falter in processing vast datasets and adapting to rapid market changes (Gamil et al., 2007; Luo et al., 2002).

This paper presents ElliottAgents, an multi-agent system that harnesses the power of NLP (Lane et al., 2019) and LLMs to analyze stock market data. Our approach combines AI-driven analysis with the Elliott Wave Principle (EWP), a established method of technical analysis (Frost et al., 2001). The core innovation lies in the system's ability to facilitate natural language dialogue

between agents, enabling them to collaboratively interpret market patterns and refine their analyses.

Our research addresses the following question: How can we effectively integrate natural language processing methods and multi-agent architectures to produce reliable and human-comprehensible stock market analyses and predictions? Through experimental validation, we demonstrate that our approach not only enhances pattern recognition accuracy but also generates detailed, easily interpretable market trend descriptions and forecasts.

Our system contributes to the field of NLP applications in specialized domains. It showcases the potential of AI-driven dialogue systems in enhancing collaborative analysis within data-intensive fields. By filling the gap between complex financial data and human understanding, ElliottAgents represents a step forward in creating more interpretable and adaptive prediction systems in finance.

## 2 Foundations of Stock Market Forecasting

### 2.1 The Evolution of Stock Market Analysis

Stock market analysis has progressed from manual techniques to AI-driven approaches over the past century. Traditional methods like fundamental analysis and technical analysis (Murphy, 1999) have been augmented by computational models since the 1960s. The development of financial software has seen several paradigm shifts: from simple automation of existing techniques to the creation of complex algorithmic trading systems (Tirea et al., 2012). Recent years have witnessed the integration of machine learning and natural language processing in financial analysis. Our proposed ElliottAgents system addresses several limitations in current market analysis approaches. The system's distributed nature enables parallel processing of market data, allowing for real-time analysis across numerous assets and timeframes simultaneously.




Figure 1: The fractal character of Elliott wave pattern (Frost et al., 2001)

In comparison to other approaches, ElliottAgents offers a more interpretable framework by integrating the structured EWP approach. Our recommendations are based on theory that has been used for years in contrast to the “blackbox” nature of other AI-based systems. This integration potentially provides a longer-term perspective and strategies more aligned with established market behavior patterns. The system’s ensemble of specialized agents, each focusing on different aspects of EWP analysis, aims to provide a more holistic market view compared to purely data-driven ensembles.

## 2.2 Elliott Wave Principle

The Elliott Wave Principle (EWP), developed by Ralph Nelson Elliott, is a technical analysis method based on the premise that market prices move in recognizable patterns driven by collective investor psychology (Frost et al., 2001; Murphy, 1999). This principle posits that market behavior alternates between phases of optimism and pessimism, creating predictable waves in price movements. Elliott identified thirteen recurring patterns, or “waves,” which can be classified into two main types: impulsive and corrective waves. As presented in Fig. 1, impulsive waves are the driving force behind market trends and consist of five sub-waves, while corrective waves, comprising three sub-waves, counterbalance the trend.

The Fibonacci sequence plays a crucial role in the EWP, providing a mathematical framework for wave relationships (Boroden, 2008). Elliott observed that waves often align with Fibonacci ratios, particularly the Golden Ratio (approximately 1.618). These ratios govern the relative lengths and




Figure 2: Fibonacci retracements in corrective waves (Frost et al., 2001)

amplitudes of waves, with Wave 3 in an impulsive sequence typically being 1.618 times the length of Wave 1. Corrective waves often retrace Fibonacci percentages (38.2%, 50%, 61.8%) of the previous impulsive wave as presented in Fig. 2.

The fractal nature of Elliott waves (Vantuch et al., 2016) allows for application across various time frames, from short-term movements to long-term trends. This characteristic, combined with Fibonacci relationships, creates a cohesive structure throughout market cycles. While the EWP does not offer certainty, it provides a framework for assessing probabilities of different market scenarios, aiding traders in understanding market context and predicting potential future paths.

## 2.3 Large Language Models

Large language models represent a significant advancement in the field of AI and NLP. These models are designed to understand, generate, and interact with human language in a way that is increasingly indistinguishable from human performance (Naveed et al., 2024; Raiaan et al., 2024). Examples of LLMs include OpenAI’s GPT-4 (OpenAI, 2023), Google’s Gemini, and the LLAMA series. The advancement of NLP is intrinsically tied to the progress of LLMs. These models, which lie at the forefront of AI, are capable of understanding, generating, and interacting with human language in a way that is increasingly indistinguishable from human performance (Louis-François Bouchard, 2024). They have been trained on vast amounts of text data and leverage sophisticated architectures to perform a wide range of language-related tasks, from translation and summarization to question answering and creative writing.

The core principle behind LLMs is the use of neural networks (NN) (Szydłowski and Chudziak, 2024b), specifically a type of network known as the transformer. Transformers have revolutionized the way models process sequential data. Unlike

traditional recurrent neural networks (RNNs) and long short-term memory networks (LSTMs), transformers can handle long-range dependencies more effectively, making them ideal for language tasks (Amaratunga, 2023). Transformers utilize a mechanism called self-attention, which allows the model to weigh the importance of different words in a sentence when making predictions. This is crucial for understanding context, as the meaning of a word often depends on the surrounding words.

While LLMs have shown potential in various NLP tasks, their application in time series prediction, particularly in finance, is an area of ongoing exploration (Tang et al., 2024). One challenge is the need for LLMs to understand the temporal order of data points (Chudziak, 2023), which is crucial for accurate forecasting (Chudziak and Cinkusz, 2024). Techniques like positional encoding are used to address this limitation (Tan et al., 2024), but further research is needed to fully leverage the capabilities of LLMs in capturing the dynamics of financial time series. The use of agents may be a factor that will greatly improve the results of time series prediction by distributing tasks among agents, enabling a more robust analysis of complex big sets of data.

### 3 Multi-Agent System Architecture

#### 3.1 System Architecture

Multi-agent systems have long been a powerful tool in modeling complex systems, where multiple autonomous entities, known as agents, interact within an environment to achieve individual or collective goals (Guo et al., 2024). Historically, these systems were built using various methodologies, including rule-based systems, symbolic equations, stochastic modeling, and early forms of machine learning.

However, these early approaches faced significant limitations. Agents were typically limited in their adaptability and often failed to respond effectively to dynamic, changing environments. Their interactions were straightforward, lacking the depth needed to mimic real-world complexities and making suboptimal decisions based on limited information and computational power.

The integration of LLMs, such as GPT-4 (OpenAI, 2023), has significantly transformed multi-agent systems, bringing advanced natural language understanding, reasoning, and decision-making capabilities to agents. LLMs enable agents to operate more autonomously, adapting to new situa-

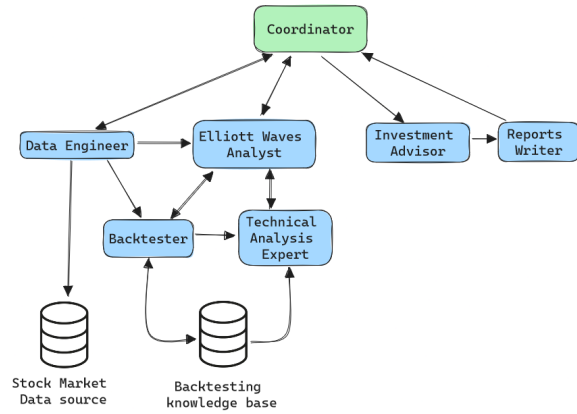


Figure 3: Data flow between agents.

tions without requiring explicit instructions. These agents can now exhibit goal-directed behaviors, making proactive decisions to achieve long-term objectives, enhancing their autonomy and proactiveness (Guo et al., 2024; Cinkusz and Chudziak, 2024). Agents can also dynamically perceive and respond to changes in their environment, learning from their experiences to improve future responses (Zhao et al., 2023; Yao et al., 2023).

The agents collaborate, performing sequential and hierarchical tasks that culminate in a comprehensive analysis as shown on Fig. 3. Some agents utilize advanced tools, which were described in section "3.2 Agents Customization".

- **Data Engineer:** The primary goal of this agent is to prepare the necessary data, that other agents will use for their analyses. This agent uses a dedicated tool that requires the name of the company, the timeframe and the interval for which the data is to be prepared, this information is provided by the user.
- **Elliott Waves Analyst:** Main task of this agent is to perform detailed Elliott waves analysis on historical stock data. To do this, we create a special tool that finds all possible impulsive and corrective wave patterns in the data. Results of this tool are used to plot charts with overlaid waves at appropriate points.
- **Backtester:** This agent uses DRL to test and validate the predictions made by Elliott waves analyst. This process allows us to identify patterns that have worked in the past on the asset that is currently analyzed, thereby we are increasing the chances of a successful analysis. The agent uses database to retain informa-

tion across past tests and can delegate tasks to other agents if necessary.

- **Technical Analysis Expert:** This agent's goal is to interpret the waves patterns with results of backtesting and choose the most likely pattern to occur in the current market state.
- **Investment Advisor:** Is responsible for synthesizing various analyses into comprehensive investment strategy. This agent uses data retrieved from RAG tool and leverages the output provided by other agents. The result of his actions is an accurate investment plan, including price levels, dates, buy or sell signal and backup plans when the future stock price does not follow the predicted trends.
- **Reports Writer:** Summarises the activities of all agents, creating a clear, easy-to-understand report for the end user. The final output is a comprehensive report that provides clear investment strategies, including when and where to buy or sell stocks, ensuring that the recommendations are up-to-date and relevant for today's market conditions.

### 3.2 Agents Customization

As presented on Fig. 4 agent is build using different components, most important technologies used by our agents are described below:

**Retrieval-Augmented Generation (RAG)** Enhances generative AI models by integrating external knowledge retrieval (Lewis et al., 2021; Asai et al., 2023). This approach converts queries into embeddings, matches them with a vectorized knowledge base, and combines retrieved data with generated responses, improving factual accuracy and reducing "hallucinations". Our system employs knowledge graph-based RAGs, which structure data into interconnected graphs (Larson and Truitt, 2024). This method improves the accuracy and relevance of generated content, allowing the model to more efficiently handle data containing a complete description of patterns and the mathematical theory behind the EWP.

**Deep Reinforcement Learning (DRL)** Combines the strengths of both deep learning and reinforcement learning (RL). It has garnered attention for its ability to solve complex problems involving sequential decision-making in high-dimensional spaces. In the traditional RL framework, an agent

learns to interact with an environment through a cycle of observing the current state, selecting an action, and receiving feedback in the form of rewards (Lapan, 2020). The agent's goal is to learn a policy, which is a strategy for selecting actions that maximizes the cumulative rewards over time. DRL enhances this process by leveraging deep neural networks, a type of machine learning model with multiple layers, to handle and approximate complex functions (Kabbani and Duman, 2022; Szydlowski and Chudziak, 2024a). Additionally, DRL can address problems with continuous action spaces, where the agent needs to select an action from an infinite set of possibilities, such as adjusting the parameters of a financial trading strategy.

DRL has been used in the backtesting process to analyze historical market data and learn effective trading strategies (Lussange et al., 2020). By identifying patterns and understanding their impact on future price movements, a DRL agent can make informed decisions to buy, sell, or hold assets, optimizing long-term returns. The ability of DRL to continuously learn and adapt proves particularly valuable in dynamic and uncertain environments, such as financial markets.

**Dynamic context** Refers to the ability of AI agents to adaptively adjust their contextual understanding based on real-time information (Witkamp, 2024). Agents can utilize various types of context, including tools, documents accessed through RAG, the history of conversations, and the ability to reflect and plan future actions. This approach leverages ongoing interactions and updates the context dynamically, enabling the agent to maintain relevance and accuracy throughout a session. By incorporating new data as it becomes available, dynamic context helps agents refine their responses and improve decision-making processes.

Memory plays a vital role in enhancing the agent's ability to understand and generate responses based on past interactions, improving decision-making and context-awareness over time. Memory in AI agents, is crucial for handling sequential data and retaining information over long periods (Weng, 2024). They achieve this through gated mechanisms that regulate the flow of information, making them highly effective for tasks requiring long-term dependencies, such as time series prediction and natural language processing.




Figure 4: Overview of a LLM autonomous agent, based on (Weng, 2024).

### 3.3 Agents Flow Engineering

Creating effective crews in multi-agent systems involves a combination of strategic orchestration, collaboration, and dynamic task decomposition (Guo et al., 2024). Multiple agents working together to achieve common goal can be managed through orchestration, where a central coordinator assigns tasks and ensures synchronized efforts. This approach enhances control and reliability by allowing the orchestrator to monitor progress, handle exceptions, and optimize resource allocation. In contrast, sequential process allows agents to interact autonomously based on predefined protocols, promoting flexibility and better decision-making.

Agents share information and work collaboratively, either through disordered cooperation (hierarchical process), where agents communicate freely and process inputs in a network-like structure, or through ordered cooperation (sequential process), where agents follow a structured sequence to build on each other’s outputs (Li et al., 2024). Our agents operates in a hierarchical mode, allowing for asynchronous execution of tasks. This significantly accelerates the entire prediction process. In this model, higher-level agents decompose complex tasks and delegate subtasks to lower-level agents. Information flows both up and down the hierarchy, with lower-level agents reporting results to their superiors, and higher-level agents providing context and coordination information to their subordinates.

Dynamic scaling is crucial for the adaptability and efficiency of multi-agent systems. By adjusting the number of active agents based on task complexity and available resources, systems can manage workloads more effectively (Guo et al., 2024). Dynamic scaling allows for the autonomous increase or decrease of agents, ensuring optimal resource utilization and maintaining system performance under varying conditions.

Another part of effective multi-agent system is task decomposition, which enables the breakdown

of complex tasks into smaller, manageable sub-tasks. In hierarchical task decomposition organizes tasks into a structured hierarchy, where each level of the hierarchy can be further decomposed until tasks reach a granularity suitable for individual agents (Chen, 2024). This clarity ensures that agents can focus on their specialized tasks while the orchestrator manages overall coordination.

## 4 Experiments and Results

### 4.1 Data and Use Cases

The presented system utilizes data from NYSE with a various intervals, allowing for the analysis of the majority of companies listed on this exchange. ElliottAgents provides the flexibility to define the time frame over which the analysis is to be performed, allowing users to conduct both short-term and long-term forecasts.

There are more patterns discovered and described in the Elliott Wave Theory and in our study we have focused on describing only a few selected ones, based on EWP we can distinguish the following use cases:

- **Identifying Impulse Waves:** Impulse waves determine the direction of the main market trend. The hypothesis is that recognizing these impulsive patterns can help predict future price movements. Impulse waves are five-wave patterns that move in the direction of the overall trend, consisting of three actionary waves (1, 3, and 5) and two corrective waves (2 and 4) (Frost et al., 2001).
- **Identifying ABC Corrections:** In EWP, ABC corrections follows the impulsive move. The hypothesis is that understanding these corrections can provide insights into potential market reversals or continuations. An ABC correction is a three-wave pattern that moves counter to the preceding impulse wave. This pattern helps traders understand when a correction is likely to end and the previous trend will resume.
- **Recognize wave extensions:** The objective of this use case is to recognize and analyze wave extensions within Elliott Wave patterns to improve prediction accuracy. Wave extensions, typically seen in the third wave of an impulsive sequence, exceed the standard 1.618 Fibonacci ratio, often reaching up to 2.618 or beyond, indicating a robust trend.

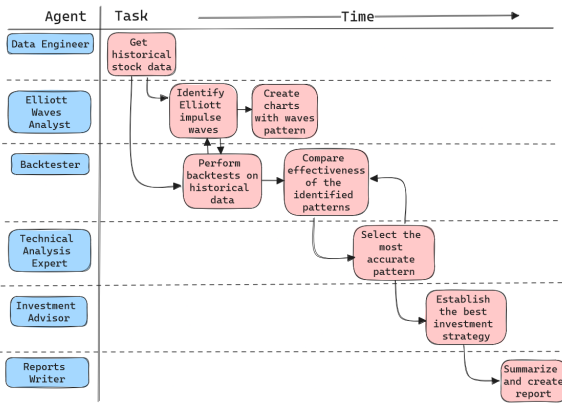


Figure 5: Flow diagram of agents identifying impulse wave use case.

- Determining support, resistance and target levels:** Support and resistance levels are critical price points where a stock is likely to reverse or pause. Support levels are price points where a downtrend is expected to halt due to a concentration of demand, while resistance levels are where an uptrend is likely to pause due to a concentration of supply. These levels are established by the ending points of previous Elliott waves.

Fig. 5 illustrates a use case diagram showing the workflow of agents in the identification of Elliott impulsive waves. Each agent is assigned specific tasks that are prerequisites for the subsequent agent's activities, ensuring a seamless and systematic process.

## 4.2 Evaluation of Use Cases

In the first part of experiment, the system tests were conducted using historical data in hourly and daily intervals. The system was run on limited historical data from the largest American companies, with data ranging from one month to two years, to recognize all waves pattern and identify possible buy or sell signals on the charts using knowledge from backtesting process. When the system issued such a signal, we iteratively added additional historical data, allowing the system to detect other patterns and issue another signals. This approach enabled us to evaluate its effectiveness in simulated, but realistic market conditions. Based on these signals, we could simulate transactions and calculate theoretical investment returns, proving the effectiveness of our agent's collaboration.

Fig. 6 presents analysis of Amazon's stock over an approximately two-month period, using hourly




Figure 6: Ending diagonal pattern recognized on AMZN 1h chart.

intervals. ElliottAgents successfully identified an "ending diagonal" pattern. This pattern, according to EWP, signifies the termination of a larger trend and often precedes a significant reversal in the market direction (Frost et al., 2001). After confirmation of the trend reversal, ElliottAgents issued a sell recommendation at \$185 per share. The target price was set at \$177 per share, which corresponds to the peak of the second wave extension within the impulsive wave sequence. As illustrated in the accompanying chart, the market behavior adhered closely to our predicted scenario. The theoretical profit from this transaction is \$8 per share, representing a 4.4% gain, achieved within a short span of just five days.

Fig. 7 presents a results of analysis conducted on Alphabet's stock over a one year period, with data aggregated on a daily interval. ElliottAgents successfully identified multiple patterns during this period. Specifically, the analysis revealed an impulsive wave sequence denoted as (1)-(2)-(3)-(4)-(5), wherein the fifth wave is an extension and a corrective wave pattern, labeled A-B-C immediately after impulsive wave. This corrective pattern terminated at the peak of the second wave of the extension, aligning perfectly with the theoretical expectations posited by EWP (Frost et al., 2001). According to the theory, the presence of this pattern suggests a forthcoming reversal exceeding the peak of the fifth wave. Upon recognizing this configuration and confirming started reversal, ElliottAgents generated a buy recommendation at a price point of \$140 per share. The target price was strategically set at \$160 per share, aligning with the peak of the fifth wave, while also accounting for the resistance level observed at the peak of wave B (\$150). This dual target strategy ensures both an optimal exit point and a buffer for potential resistance encounters. As the experimental data, in the chart shows, price levels have been achieved. The theoretical profit realized from this transaction amounted to \$20 per share, translating to a 13.3% gain.

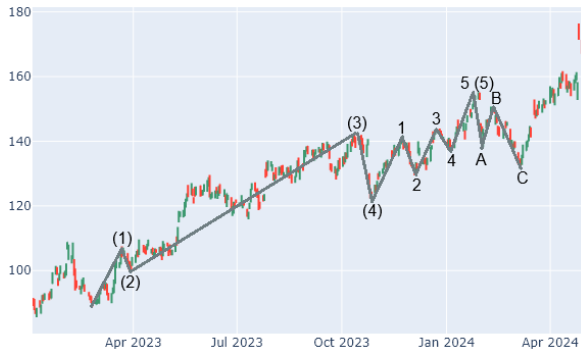


Figure 7: Fifth wave extension found on GOOG 1d chart.

In the Fig. 8 we present a long-term analysis of Nvidia’s stock price, on daily intervals. System detected a complete Elliott wave cycle, consisting of an impulsive wave followed by a corrective wave. The identification of this pattern suggests the potential end of the corrective phase and the continuation of the broader trend (Murphy, 1999), which is bullish in this case. The peak of wave C was precisely identified at \$39 per share. Anticipating a reversal in the trend, the system issued a buy recommendation at \$42 per share. The target price was set at \$50 per share, located at the peak of the fifth wave. This target signifies the emergence of the first wave in a new impulsive sequence, according to EWP. The chart clearly demonstrates that a sharp rebound occurred shortly after the recommendation was made, resulting in the target price being reached. This scenario yielded a theoretical profit of \$8 per share, representing a 17.4% gain.

The second part of the experiment focus on quantity tests for the correctness of the detected pattern and the impact of DRL on results. Test was conducted using a cross-validation method on 1000 samples (candlesticks) with a daily interval for the



Figure 8: Full wave cycle recognized on NVDA 1d chart.

stocks, which were analyzed in the previous section. In these tests, we focused on examining unfinished impulsive waves (1-2-3-4) consisting of 4 sub-waves as well as complete impulsive waves (1-2-3-4-5), in each case waves could not overlap. In both cases, we compared the results with and without a DRL backtesting process conducted on 10 years of historical data for each company.

Based on the identified patterns, agents predicted whether the next movement would be upward or downward. A prediction was considered correct if the average price of the subsequent  $n$  candlesticks was higher or lower, depending on the issued signal. The  $n$  number of candlesticks was determined according to EWP, where in the case of waves 1-2-3-4, the length of the fifth wave should be approximately 1.62 times the length of the first wave, and in the case of a complete impulse wave, the following wave A should have a length close to wave 5.

Table 1 presents the results of the cross-validation experiments for 1000 data samples in two time intervals. As we can see, the identification of a complete impulsive wave pattern contributes to better predictions of subsequent price movements than incomplete impulse wave pattern. In case of hourly intervals our system detected smaller number of patterns, mainly because price changes on the hourly interval were smaller. The use of DRL resulted in an improvement in prediction, showing that agents are able to use the learning process on historical data in better interpretation of patterns.

### 4.3 Success Criteria

The success criteria for ElliottAgents focus on accurate pattern recognition and analysis to enable users to achieve real market profits. By leveraging NLP, the system can present complex financial information in a clear and understandable manner. The system must provide analysis of a company’s stock based on user inputs, identifying all possible wave patterns. These patterns should be visually represented on a chart, and based on them, agents should provide actionable insights, including target price levels based on Fibonacci relationships and key support and resistance levels. Additionally, it should offer clear buy or sell recommendations based on wave analysis, price projections, and trend analysis, with specific time frames for action. Success is measured by the system’s ability to deliver pattern recognition and analysis that can be used by traders in investment decisions.

Stock	1-2-3-4 Patterns			1-2-3-4-5 Patterns		
	N	Without backtesting	With backtesting	N	Without backtesting	With backtesting
<i>Daily Interval</i>						
AMZN	24	58.34%	66.67%	18	66.67%	77.78%
GOOG	28	53.57%	67.86%	23	65.22%	82.61%
INTC	19	57.89%	73.68%	15	60.00%	73.34%
<i>Hourly Interval</i>						
AMZN	10	50.00%	70.00%	8	62.50%	75.00%
GOOG	13	53.84%	61.54%	9	77.78%	77.78%
INTC	12	58.34%	66.67%	9	66.67%	88.89%

N: number of patterns found.

Table 1: Comparison of pattern recognition with and without backtesting

## 5 Discussion and Future Work

### 5.1 Comparison with Other Systems

Multi-agent architectures have been utilized in stock price prediction systems for many years (Akintola and Oyetunji, 2021; Gamil et al., 2007; Luo et al., 2002). However, advancements in AI over recent years have significantly enhanced these systems capabilities. Traditional systems often relied on static rules and fuzzy logic to make decisions, but they faced limitations in accuracy and adaptability. The introduction of fuzzy logic, as seen in older systems, provided a foundation for integrating qualitative judgments with quantitative analysis, yet it required further optimization to improve decision-making. It is difficult to compare the profitability of our system with other price prediction systems available to date. However, based on our experiments, we see that the system can effectively detect and interpret wave patterns, with better accuracy than similar systems using EWP (Tirea et al., 2012). The analyses created by our agents, clearly present an investment plan, with price levels, that can be used in real world by the traders.

### 5.2 Future Enhancements

Currently, our work has focused primarily on a few patterns recognized by EWP. Expanding our analysis to include additional wave formations such as truncations, zigzags, flat corrections, triangles, and other patterns could significantly enhance our predictive capabilities. Following the successful integration of EWP, we could further improve our system by incorporating other technical analysis methods, such as moving averages. This expansion could enhance our ability to determine more

accurate buy or sell signals, potentially improving signal reliability and profitability.

## 6 Conclusion

ElliottAgents demonstrates the potential of integrating NLP and multi-agent systems in the domain of stock market analysis (Tunstall et al., 2022). By leveraging LLMs and the EWP, the system transforms complex historical market data into comprehensible predictions and explanations. The key innovation lies in the inter-agent dialogue, which mimics collaborative human analysis while harnessing AI’s pattern recognition capabilities. This approach not only enhances the accuracy of technical analysis but also addresses the challenge of making financial data interpretable to human users.

Experimental results, conducted on historical data over a period of several years on some of the largest U.S. companies, validate the system’s effectiveness in recognizing market patterns and generating natural language descriptions of trends across various time frames. The multi-agent architecture, facilitated by advanced NLP techniques, enables the decomposition of complex analytical tasks, leading to more nuanced and reliable predictions. This research contributes to the broader field of NLP applications in data-intensive domains, showcasing how AI-driven dialogue systems can enhance collaborative analysis. ElliottAgents bridges the gap between sophisticated AI analysis and human understanding, paving the way for more interpretable and adaptive prediction systems in finance and potentially other specialized fields.



## References

- K.G. Akintola and O.E. Oyetunji. 2021. Development of an agent-based framework for stock market trading. *IRE Journals*, 4(9).
- Thimira Amaratunga. 2023. *Understanding Large Language Models Learning Their Underlying Concepts and Technologies*. apress.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. Self-rag: Learning to retrieve, generate, and critique through self-reflection. arXiv:2310.11511 [cs.CL].
- Carolyn Boroden. 2008. *Fibonacci Trading: How to Master the Time and Price Advantage*. McGraw Hill.
- Haiping Chen. 2024. Understand the llm agent orchestration. <https://medium.com/scisharp/understand-the-llm-agent-orchestration-043ebfaead1f>. Accessed: Jun. 1, 2024.
- Adam Chudziak. 2023. Predictability of stock returns using neural networks: Elusive in the long term. *Expert Systems with Applications*, 213.
- Jaroslav A. Chudziak and Konrad Cinkusz. 2024. Towards llm-augmented multiagent systems for agile software engineering. In *The 39th IEEE/ACM International Conference on Automated Software Engineering (ASE 2024)*, Sacramento, CA, USA.
- Konrad Cinkusz and Jaroslav A. Chudziak. 2024. Communicative agents for software project management and system development. In *Proceedings of the 21th International Conference on Modeling Decisions for Artificial Intelligence (MDAI 2024)*, Tokyo, Japan.
- A. J. Frost, Robert R. Prechter Jr., and Charles J. Collins. 2001. *Elliott Wave Principle: Key to Market Behavior*. Wiley.
- Ahmed A. Gamil, Raafat S. El-fouly, and Nevin M. Darwish. 2007. Stock technical analysis using multi agent and fuzzy logic. In *Proceedings of the World Congress on Engineering, WCE 2007*, London, UK.
- Taicheng Guo et al. 2024. Large language model based multi-agents: A survey of progress and challenges. arXiv:2402.01680v2 [cs.CL].
- Taylan Kabbani and Ekrem Duman. 2022. Deep reinforcement learning approach for trading automation in the stock market. *IEEE Access*, 10.
- Hobson Lane, Hannes Hapke, and Cole Howard. 2019. *Natural Language Processing in Action: Understanding, analyzing, and generating text with Python*. Manning.
- Maxim Lapan. 2020. *Deep Reinforcement Learning Hands-On Second Edition*. Packt.
- Jonathan Larson and Steven Truitt. 2024. Graphrag: Unlocking llm discovery on narrative private data. <https://www.microsoft.com/en-us/research/blog/graphrag-unlocking-llm-discovery-on-narrative-private-data>. Accessed: May. 10, 2024.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2021. Retrieval-augmented generation for knowledge-intensive nlp tasks. arXiv:2005.11401v4 [cs.CL].
- Junyou Li, Qin Zhang, Yangbin Yu, Qiang Fu, and Deheng Ye. 2024. More agents is all you need. arXiv:2402.05120v1 [cs.CL].
- Louie Peters Louis-François Bouchard. 2024. *Building LLMs for Production: Enhancing LLM Abilities and Reliability with Prompting, Fine-Tuning, and RAG*. Towards AI.
- Y. Luo, Kecheng Liu, and Darryl N. Davis. 2002. A multi-agent decision support system for stock trading. *IEEE Network*, 16(1).
- Johann Lussange, Ivan Lazarevich, Sacha Bourgeois-Gironde, Stefano Palminteri, and Boris Gutkin. 2020. [Modelling stock markets by multi-agent reinforcement learning](#). *Computational Economics*. Hal-03055070.
- John J. Murphy. 1999. *Technical Analysis of the Financial Markets: A Comprehensive Guide to Trading Methods and Applications*. New York Institute of Finance.
- Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Muhammad Usman Saeed Anwar, Naveed Akhtar, Nick Barnes, and Ajmal Mian. 2024. A comprehensive overview of large language models. arXiv:2307.06435 [cs.CL].
- OpenAI. 2023. Gpt-4 technical report.
- Mohaimenul Azam Khan Raiaan, Md. Saddam Hossain Mukta, Kaniz Fatema, Nur Mohammad Fahad, Sadman Sakib, Most Marufatul Jannat Mim, Jubaer Ahmad, Mohammed Eunus Ali, and Sami Azam. 2024. A review on large language models: Architectures, applications, taxonomies, open issues and challenges. *IEEE Access*, 12.
- Kamil L. Szydlowski and Jaroslav A. Chudziak. 2024a. Toward predictive stock trading with hidformer integrated into reinforcement learning strategy. In *The 36th International Conference on Tools for Artificial Intelligence (ICTAI 2024)*, Herndon, VA, USA.
- Kamil L. Szydlowski and Jaroslav A. Chudziak. 2024b. Transformer-style neural network in stock price forecasting. In *The 21th International Conference on Modeling Decisions for Artificial Intelligence (MDAI 2024)*, Tokyo, Japan.

- Mingtian Tan, Mike A. Merrill, Vinayak Gupta, Tim Althoff, and Thomas Hartvigsen. 2024. Are language models actually useful for time series forecasting? arXiv:2406.16964v1 [cs.LG].
- Hua Tang, Chong Zhang, Mingyu Jin, Qinkai Yu, Zhen-ting Wang, Xiaobo Jin, Yongfeng Zhang, and Mengnan Du. 2024. Time series forecasting with llms: Understanding and enhancing model capabilities. arXiv:2402.10835v2 [cs.CL].
- Monica Tirea, Ioan Tandau, and Viorel Negru. 2012. Stock market multi-agent recommendation system based on the elliott wave principle. In *International Conference on Availability, Reliability, and Security*, Prague, Czech Republic.
- Lewis Tunstall, Leandro von Werra, and Thomas Wolf. 2022. *Natural Language Processing with Transformers, Revised Edition*. O'Reilly.
- Tomas Vantuch, Ivan Zelinka, and Pandian Vasant. 2016. Market prices trend forecasting supported by elliott wave's theory. In *First EAI International Conference on Computer Science and Engineering*, Penang, Malaysia.
- Lilian Weng. 2024. Llm powered autonomous agents. <https://lilianweng.github.io/posts/2023-06-23-agent/>. Accessed: Jun. 1, 2024.
- Frank Wittkamp. 2024. Next-level agents: Unlocking the power of dynamic context. <https://towardsdatascience.com/next-level-agents-unlocking-the-power-of-dynamic-context-68b8647ee689>. Accessed: Jun. 1, 2024.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. React: Synergizing reasoning and acting in language models. arXiv:2210.03629 [cs.CL].
- Pengyu Zhao, Zijian Jin, and Ning Cheng. 2023. An in-depth survey of large language model-based artificial intelligence agents. arXiv:2309.14365v1 [cs.CL].