



Tradesense - AI-Powered Stock Analysis and Trade Stimulation Platform

Vansh¹, Manvi², Abhishek Gulia³

Department of Computer Science & Engineering (AI & DS), Panipat Institute of Engineering and Technology, Panipat, India^{1,2,3}

vanshbhardwajhere@gmail.com¹, manvisharma2435@gmail.com², abhishek.cse.ai-ds@piet.co.in³

Abstract. *Financial markets have experienced a rapid development and the increasing involvement of retail investors has resulted in a high demand of smart and user-friendly decision support systems. The conventional tools of financial analysis generally demand substantial expertise in the domain, and do not offer user-friendly interaction to non-expert users. To solve this problem, this paper proposes TradeSense, an AI-based financial analytics system that combines Large Language Models (LLMs), real-time financial information APIs, and cloud-based computing to allow the analysis of investments based on natural language input. The system enables users to type in plain language financial queries which are executed by deploying a high-performance LLM through the Groq API to understand user intent and produce formatted analytical queries. The suggested solution gathers real-time financial information through the Yahoo Finance API, which was chosen due to its flexibility and limited rates, and predictive analytics to determine possible investment performance, based on the past. The TradeSense is powered by a scalable architecture that uses Firebase as an authentication server, database, and frontend, and Hugging Face as the backend server. Also, the system keeps a history of user queries so that it can be used more easily and customized. The combination of natural language processing, real-time data retrieval, and predictive modeling within a single system makes TradeSense stand out of the current solutions and can help make financial analytics more approachable, interactive, and efficient to a broad spectrum of users.*

Keywords: TradeSense, Financial Analytics, Stock Market Prediction, Large Language Models (LLMs), Natural Language Processing (NLP), Groq API, Yahoo Finance API, Firebase, Cloud Computing, Machine Learning, Investment Analysis, Decision Support System.

Introduction

The hyper growth of the world financial markets, as well as the increasing involvement of the retail investors, has dramatically changed the environment of investment decision making. Due to the presence of online trading sites and online financial tools, people have never had much access to market information as they do today. But the dynamic and complicated nature of financial markets makes it to be difficult to process information effectively and make sound judgment by non-expert users. Conventional FY tools may need a lot of domain familiarity, technical skills and understanding of analytical indicators and hence they cannot be used by a wider audience. Over the past several years, new opportunities in simplifying financial analysis have been made with the development of artificial intelligence, specifically in the field of machine learning and natural language processing. Large Language Models (LLMs) have been shown to comprehend and process human language, which can then be used to create intelligent systems that can



enable users to work with complex data by using natural language queries. Regardless of these developments, most of the current systems are either entirely based on predictive modelling or they offer static financial information without an integrated mechanism that incorporates real time data access and smart query processing along with easy user interface.

In a bid to overcome these shortcomings in this paper the author introduces TradeSense, an artificial intelligence-based financial analytics and decision support tool that is aimed at filling the void between complex financial data and accessibility to users. The system allows natural language and can be used to enter investment-related queries, which are handled by a high-performance LLM through the Groq API. The resulting interpreted queries are then adopted to get real-time financial information on the Yahoo Finance API to ensure that the information is current and reliable. The system also includes predictive analytics to forecast the possible investment results, giving the users a better understanding of the trends in the market. TradeSense is designed on a scalable cloud framework based on Firebase as an authentication service and database services as well as frontend services are hosted on Firebase and backend services are deployed through Hugging Face. The system is also quick to keep a history of user queries to provide enhanced personalization and usability. TradeSense provides a new model of financial analytics, integrating natural language processing, real-time data discovery, and predictive modeling into one platform that is easy to access and efficient. The main outputs of the work are the creation of an LLM-based system of financial queries, the incorporation of real-time financial APIs and predictive analytics, and the deployment of a fully operational, scalable cloud-based system. The suggested system will be a democratization of financial analysis as it will allow users of all levels of technical skills to make informed investment choices using a user-friendly and interactive interface.

2. Related Work

Financial analytics and stock market prediction is an area that has been experiencing high development due to the use of machine learning, technical analysis and natural language processing. Other studies have been done to enhance predictability, data availability and ease of interaction to the users. The majority of the existing systems however confront such aspects individually instead of offering an integrated solution. In this section, the review of the pertinent work in key areas relating to the proposed TradeSense system is done.

2.1 Machine Learning-Based Stock Prediction

Machine learning has widely been applied to predict stock market. Patel et al. (2015) discussed the algorithms like Support Vector machine (SVM), random forest, and artificially neural network to forecast stock price changes using historical data [1]. Their findings revealed that machine learning models have the capability of identifying patterns in financial data, but that they tend to be affected by market volatility and do not extrapolate well to unobservable conditions. Likewise, Zhang et al. (2020) offered deep learning based on Long Short-Term Memory (LSTM) networks to predict time-series [2]. The model proved to have better accuracy as it was able to capture temporary stock data correlations however it needed huge volumes of data and needed large-scale computational resources hence could not easily be implemented in real time.

2.2 Technical Analysis-Based Systems

A number of studies have been devoted to the analysis of financial data with the help of technical indicators like Moving Averages, Relative Strength Index (RSI), and Moving Average Convergence Divergence (MACD). Kim and Han (2016) came up with a system that used these signs to determine



trends and come up with trading signs [3]. Although this type of system would help in the analysis of historical data, it is not adaptive as it is based on the existing rules. Furthermore, they demand users to read complicated charts by hand, which is not as appropriate with the unprofessional users.

2.3 NLP and LLM-Based Financial Assistants

Advances in Natural Language Processing (NLP) have seen more recent systems implemented with conversational interfaces to enhance accessibility [10],[12]. Large Language Models (LLMs) allow users to query financial systems through natural language. These systems help improve the user experience since they decrease the technical knowledge. Nonetheless, the vast majority of the available implementations offer only fixed or partial answers and do not work well with real-time financial data or predictive analytics allowing the researcher to limit the overall utility of these solutions in dynamic markets.

2.4 Financial Data Retrieval Using APIs

One of the most important parts of the contemporary analytics is financial data retrieval. Market data is normally accessed via APIs like Alpha Vantage and Yahoo Finance [5], [6]. Although Alpha Vantage is well used in research, it has a very tight rate limit, making it difficult to use in real-time. In comparison, Yahoo Finance is the platform that gives users more flexibility and reliability in accessing financial information. Nevertheless, there are numerous systems which are based on fixed datasets, as opposed to dynamic data retrieval based on API, which makes them less relevant to real world scenarios.

2.5 Cloud-Based Financial Systems

Financial applications have been deployed in a scalable and efficient manner through cloud computing. Authentication, database management, and hosting services are offered as a bundle on platforms like Firebase, and frameworks like Hugging Face allow scalability of the backend [7], [8]. Nevertheless, numerous studies are theoretical and lack illustration of end-to-end systems. They do not have cohesive user interfaces, processing and data storage units.

2.6 Research Gap and Motivation

Although there has been a lot of advancement in the individual fields of machine learning, technical analysis, NLP, and cloud computing, the current systems do not integrate these aspects. It lacks solutions that are integrated and combine natural language understanding, real-time financial data retrieval, predictive analytics, and scalable cloud infrastructure. As a solution to this, the proposed TradeSense system adds a Large Language Model through the Groq API to process natural language queries, Yahoo Finance API to retrieve real-time data, and Firebase as well as Hugging Face to scale the implementation. This centralized solution makes it possible to have a smart, interactive, and easy-to-use financial analytics platform.

3. Proposed Solution

The TradeSense system is proposed as an AI-based financial analytics system that combines job language processing, real-time data retrieving, predictive analytics, and cloud-based implementation. The system allows users to query financial data with natural language queries and get valuable insights in an intuitive way. Workflow of the entire system is further broken down into various stages with each overseeing a particular functionality.

3.1 User Authentication Module

Firebase Authentication is used in the system to provide smooth user access in a safe and secure manner [7]. The user can use Google Sign-In to log in, a feature that has made the authentication process easier and



improved user experience. In this module, only the authorized users can access the system and can customize it by keeping user specific data.

3.2 Natural Language Query Processing

Users are then authenticated and are able to type financial queries in natural language. These queries are run through a Large Language Model (LLM) combined through the Groq API [9]. The LLM is able to understand the intent of the user and break down the unstructured text into a structured format that can then undergo further processing. This allows users access to the system without the need of technical expertise in financial analysis.

3.3 Financial Data Retrieval Module

Once query is processed, financial data is retrieved based on the query through Yahoo finance API [5]. This API was chosen as opposed to Alpha Vantage because it is more flexible and has less rate limits than it thus lends more to real time applications [6]. The module retrieves both past and present financial information which is needed to conduct analysis/calculations.

3.4 Data Processing and Analysis

The data that is retrieved is then used to calculate the applicable financial indicators like the returns on investments, movements in price and trends. The system implements analytical methods and can use technical indicators to increase the quality of insights. It is a process of converting raw data into information which is easily interpretable by the users [1].

3.5 Prediction Module

The system uses a machine learning-based prediction module in order to offer a futuristic look into the system. The module uses historical data to predict the possible values of the assets in the future. The system will augment the decision-making skills and give the users a more accurate perspective of the potential returns on investments by incorporating forecast analytics.

3.6 Visualization Module

The results are presented in the form of interactive user interface. The information is more visible, and easier to comprehend by visualizing the trends, price moving, and growth of investment using graphs and charts. This improves the interaction with the users and decision making.

3.7 Data Storage and History Management

Firebase firestore is used to store queries and other pertinent results to the user in the application [7]. This enables you to save a history of queries in order to reconnect with Marty. It also is good to customization and implementation.

3.8 System Workflow

The general operation of the system is as follows: A user authenticates with Firebase Authentication [7] to enter his/her data, and this is followed by a natural language query that is processed through the Groq LLM [9] and the information is returned to Yahoo Finance. The retrieved data is then analyzed and possibly processed by the prediction module [2] whereby the results are then visualized and stored in the database to be used.

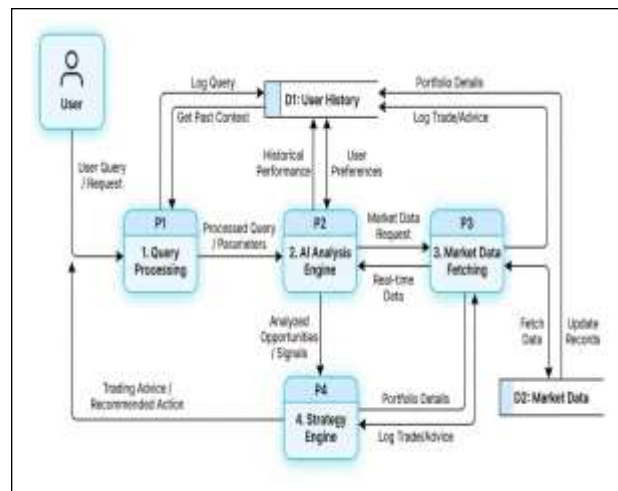


Fig. 1: System Workflow.

The combination of these modules will lead to a coherent, scaled-up, and smart financial analytics solution that tackles the flaws of current strategies and makes it easier to use in practice.

4. System Architecture

The system is developed based on a layered and modular architecture which structures various functional components in separate units. This design of the architecture will ensure the ability to scale, maintainability and communication efficiencies between the system components. Unlike methodology that is focused on step-by-step workflow, the system architecture is more focused amid structuring of the system and in its relation to what modules it consists of.

The architecture is further designed into different layers and each segments has its own specific set of operations. Each of these layers interacts with another by making data flow and user query processing easy and real time respectively.

4.1 Presentation Layer

The front face of the system is the Presentation Layer where the end-users have the opportunity to interface with the application. It provides customer interface of entering the financial questions and provides response. It is this layer which can produce the visual output e.g. charts, graphs and textual information to enable user experience easy and user friendly user experience is attained.

4.2 Authentication Layer

Authentication Layer consists of user authentication and Firebase Authentication [7] is employed to provide access. It enables Google Sign-In to provide a sensible and secure way of logging in. This level will ensure that the system operates solely according to who is authorized to do so and it will enable personal control of data.

4.3 Application Layer (LLM Processing)

The important processing unit of the system is the Application Layer. It is a combination of a Large Language Model (LLM) and Groq API to respond to user-queries [9]. It is a layer endowed with the role of trying to interpret the purpose of the user and generating structure representations as required to execute the next operations [10].



4.4 Data Integration Layer

The Data Integration Layer deals with the communication with the external financial sources. It uses the Yahoo Finance API to get real-time and historical financial data [5]. It is an API layer that sifts out the API contacts and offers accessible and stable information access [6].

4.5 Analytics and Prediction Layer

This layer fulfills the functions of analytical calculations and prediction. It deciphers the recovered information to come up with financial inferences and makes machine-learning predictions about the future trends [2], [11]. This decoupling of this layer ensures modularity and scalability of the system.

4.6 Data Storage Layer

The Data Storage Layer uses Firebase Firestore to store user questions, answers and history of communication. It also ensures efficient data management and supports such features as the history of tracks query and personalization [7].

4.7 Deployment and Hosting Layer

Deployment Layer deals with model execution and hosting backend services. It has the Hugging Face to launch machine learning models and to process back-end tasks [8]. The layer ensures that it is scaled, reliable and is an efficient system.

Table I: System Architecture Components

Layer / Module	Description
Presentation Layer	Provides user interface for entering queries and displaying results through charts and text
Authentication Layer	Manages secure login using Firebase Authentication (Google Sign-In)
Application Layer (LLM)	Processes natural language queries using Groq API and converts them into structured format
Data Integration Layer	Fetches real-time and historical financial data using Yahoo Finance API
Analytics & Prediction Layer	Performs financial calculations and applies machine learning models for forecasting
Data Storage Layer	Stores user queries and results using Firebase Firestore for history tracking
Deployment Layer	Hosts backend services and models using Hugging Face for scalability and performance

The multi-layered architecture feature of the TradeSense design allows it to separate parts of the software allowing individual components to function autonomously, but aid in the collective functionality of the system. This design makes the system highly maintainable, it enhances the performance and provides future scalability and ability to add a new feature.

5. Result and Discussions

A proposed system TradeSense was implemented and tested to examine its effectiveness of delivering smart financial information using natural language queries. The system involves the integration of several



modules, such as authentication, query processing, data retrieval, prediction, and visualization into one platform. The findings prove the usefulness of integration of Large Language Models with real-time financial data and cloud-based infrastructure.

5.1 Query Interpretation Performance

The system was put to test with different user queries that pertained to investment returns, asset comparison, time-based analysis. It was noted that the Large Language Model implemented using the Groq framework is able to process natural language inputs and then translates them into structured queries. The responses generated were correct and corresponded to the intention of the user, which proves the efficiency of processing with LLM.

5.2 Real-Time Data Retrieval Efficiency

Yahoo Finance API integration helped to perform efficient access to the real-time and historical financial data. The system was more reliable and had low latency when compared to other APIs like Alpha Vantage because of the reduced rate limits. This guaranteed the efficiency of performance of the system in terms of executing several queries.

5.3 Prediction Module Evaluation

The prediction module was tested depending on how it predicts the trends in the future using past data. The findings reveal that this model gives reasonably good estimates, which help the users to know about possible future results. The predictions can serve as good directional clues although financial markets are volatile in nature.

5.4 Visualization and User Experience

The system improves the knowledge of the users because it displays the results in form of line charts and trend graphs. This enhances interpretability and lets the financial data be easily analysed by the user. Easy to understand interface makes it easily available to the technical and non-technical user.

5.5 Data Storage and History Tracking

Firebase Firestore is an effective data structure that maintains the user query and its result. This allows having the history of queries and allows returning to the repeated analyses. The storage system enhances individualization and general user experience.

5.6 Comparative Analysis with Existing Systems

Compared to the current systems, TradeSense has a number of benefits, such as the opportunity to interact using the natural language, getting the real-time access to the data, being deployed on the cloud, and possessing the integrated predictive functionality. All these make the system more efficient and easier to use than the tradition financial tool as shown in fig. 2.



Fig. 2: Features Comparison.



5.7 Limitations of the System

The system has its limitations in spite of its merits. It requires external APIs to retrieve data, this can result into performance problems in case of failure of network or API. Also, the effectiveness of predictions is determined by the quality and accessibility of past data.

5.8 Overall Discussion

The general findings reveal that TradeSense provides a successful way of bridging the gap between the complex financial analysis and the level of user-friendliness. With sophisticated technologies incorporated into one system, the system will offer feasible, scalable, and smart financial decision-making.

Table II: Performance Evaluation Metrics

Metric	Description
Query Accuracy	Measures how correctly the system interprets user queries
Response Time	Time taken to process and return results
Prediction Accuracy	Accuracy of forecasted values compared to actual data
API Latency	Time taken to fetch data from Yahoo Finance
User Satisfaction	Based on usability and ease of interaction

6. Conclusion and Future work

The present paper introduced TradeSense, an artificial intelligence-based Stock Analysis and Trade Stimulation system, a combination of natural language processing, real-time data search, predictive traffic simulation, and the cloud-based technology. The system allows users to communicate with financial information through the natural language queries, without having to be technical in market analysis. The system manages to effectively decode user intent and translate it into structured analytical operations by using a Large Language Model via the Groq API [9], [10]. By integrating the Yahoo Finance API, you can have access to the real-time and historical financial information [5], and the prediction module will produce additional power to the system by offering you the information about the direction of financial trends in the future. Firebase Authentication, database management, and hosting, as well as Hostginger Face backend deployment, guarantees scalability and effective system performance. The findings indicate that proposed system offers real-time, convenient and precise financial information and hence is a viable tool to both novice and seasoned investors. Regardless of the benefits, the system has some limitations. It can also access the external APIs to access the data, which may also serve as a performance hitched issue when the network goes down, or the API is not functioning properly [5], [6]. In addition, the predictive ability will depend on the quality of past data and model of the machine learning taken. Financial markets are generally volatile and these further declines prediction outputs accuracy.



Moving the work forward, however, can only be done in the future by creating a more sophisticated deep learning model that will yield a better prediction result. Another factor that could amplify the decision-making capacities of the data gathering is sentiment analysis (combined with news and social media data). Moreover, the system can be adjusted to include the capabilities of the portfolio management, automated trading and maintenance of the multiple asset classes. It can also improve performance of the system by ensuring that the system is made more reliable, through better caching techniques and reduced dependence on external APIs.

TradeSense in general seems to provide a smart, promising scaleable and user-friendly platform, where the intricate financial operations analytics can be made available in the hands of users in the creation of AI-driven financial technology.

References:

- [1] H. Patel, S. Shah, P. Thakkar and K. Kotecha, Predicting stock market index using fusion of machine learning techniques in *Expert Systems with Applications*, vol. 42, no. 4, pp. 2162–2172, 2015.
- [2] Y. Zhang, Y. Aggarwal and S. Qi, Stock price prediction by LSTM neural network, in *Proceedings of the IEEE International Conference on Big Data*, 2020, pp. 4423–4427.
- [3] K. Kim and I. Han, Genetic algorithms approach to feature discretization in artificial neural networks as the predictor of stock price index, *Expert Systems with Applications*, vol. 19, no. 2, pp. 125–132, 2016.
- [4] J. Brownlee, *Deep Learning Time Series Forecasting, Machine Learning Mastery*, 2018.
- [5] Yahoo Finance API, 2019, Yahoo Finance Market Data, [Online]. Available: <https://finance.yahoo.com/>. [Accessed: Mar. 2026].
- [6] Alpha Vantage API, Realtime and Historical Stock Data API, [Online]. Available: <https://www.alphavantage.co/>. [Accessed: Mar. 2026].
- [7] Firebase, “Firebase Documentation, Google, [Online]. Available: <https://firebase.google.com/>. [Accessed: Mar. 2026].
- [8] Hugging Face, “Machine Learning Model Deployment Platform, [Online]. Available: <https://huggingface.co/>. [Accessed: Mar. 2026].
- [9] Groq, “Groq API Documentation to LLM Inference, [Online]. Available: <https://groq.com/>. [Accessed: Mar. 2026].
- [10] T. Mikolov et al., “Efficient estimation of word representations in vector space, in *Proceedings of Workshop at ICLR*, 2013.
- [11] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [12] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, MIT Press, 2016.