Nexus Gate Fund

Najib Hersi
California State University, Fullerton
800 North State College Boulevard
Fullerton, California
nhersi1@csu.fullerton.edu

Sebastien Minassian
California State University, Fullerton
800 North State College Boulevard
Fullerton, California
sebastienminassian@csu.fullerton.edu

Joshua Castaneda
California State University, Fullerton
800 North State College Boulevard
Fullerton, California
jcastaneda85@csu.fullerton.edu

Haley Patel
California State University, Fullerton
800 North State College Boulevard
Fullerton, California
haleyjpatel@csu.fullerton.edu

ABSTRACT

Many university students struggle to compete in the job market and pay for necessities, ultimately leading them to struggle to have money. University students need to have another way to earn money while also having time for their academic career. Many options to invest are not accessible to university students or require a lot of time set aside to continuously learn about the market. This paper presents Nexus Gate Fund, a stock trading bot that makes trading decisions on behalf of the user, meaning that the user does not have to be heavily involved in the trading process. The algorithm that Nexus Gate Fund uses involves checking the sentiment of recently published news articles to determine if it should buy, hold, or sell a stock. The algorithm is currently still being developed and refined for the past six months to increase the portfolio revenue. Results have shown that the refinements have increased the portfolio revenue. We will continue to develop and refine the algorithm to increase the revenue the bot makes.

CCS Concepts

• Software and its engineering → Software creation and management → Designing software → Software design engineering

Keywords

Software; Engineering; Creation; Management; Designing;

1. INTRODUCTION

Money is a crucial aspect of people's lives, and it is needed to get resources such as shelter and food. However, young people are struggling to earn it. Many university students struggle with finding a job in a highly competitive market. In addition, the rising costs of food and rent make many students worry about their financial situation. Although investing can allow them to generate wealth without spending time, there are many factors that overwhelm and prevent college students from doing so.

One factor that may prevent university students from investing in the stock market is the lack of financial literacy. Despite the rising debt that university students have, many students have little knowledge regarding financial literacy. A recent study shows students that were tested on their financial literacy; only 14.5% of students answered all six questions correctly [8]. The lack of financial literacy can make students struggle to pay back their debts and not utilize their money as efficiently as they can. Students need to set aside time in their schedules to learn how to become financially literate.

Furthermore, students also need to make time to learn about the complexities of investing before jumping right in. It is advised to research what to use to invest and initial stocks to purchase. To add on, people need to keep up to date with news regarding companies to know when to buy, hold, or sell stocks. Due to university students' hectic schedules, they may not spend time learning how to invest in the stock market.

In this work, we propose an automated trading bot that can buy, hold, and sell stocks on behalf of users. The following is a summary of the major technical contributions made:

1) The trading bot can successfully do actions on behalf of the user and perform trading decisions: buy, hold, or sell the stock. These decisions are based on the sentiment of the latest market news. The bot is currently using paper money to trade.
2) The decisions of the trading bot are displayed to the user in the Portfolio section. Users can see statistics such as the current value of portfolio and return percentage.
3) A "Market News" tab is available for users to click and view the top market news. Users can get an understanding of what events are occurring that can cause the price of stocks to differ.

The rest of the sections in the article are structured as follows. In Section 2, we will elaborate on the relevant background. In Section 3, we will present the proposed Nexus Gate Fund, an automated trading bot. The results of the automated trading bot will be presented in Section 5. Finally, Section 6 will conclude the article.

2. RELATED WORK

2.1 Quantitative Hedge Funds

Quantitative Hedge Funds like Citadel, Two Sigma, and Renaissance Technologies are some of the most advanced trading systems in the world. These companies use highly trained teams of researchers with backgrounds in science, math, and physics to build sophisticated statistical models that allow for high frequency trading in microseconds speeds. Renaissance Technologies' Medallions Fund has had an average return on investment (ROI) prior to fees of 66% per year over a 30-year period [3]. Average annual ROI prior to fees from 1994 through mid 2014 was 71.8[4].

These top-tier trading strategies are only available to an elite number of investors who are current or former employees of Medallion Funds, and their immediate family members [5]. Access to the Medallion Fund has been closed off to outside investors since 1993.

Furthermore, to invest in hedge funds, an investor must qualify as an accredited investor. This means that the investor will have an annual income of at least $200,000 for individuals ($300,000 for couples) for the previous two years [6]. The average university student, who is often managing student loan debt and working a job, does not have the means of obtaining these trading tools.

2.2 Robo-advisors

With platforms like Wealthfront and Betterment, Robo-advisors have created automated portfolio management options for non-accredited, regular investors. These platforms utilize an algorithm-based approach to allocate their resources. As a result, robo-advisors create diverse asset allocations within low-cost exchange-traded funds (ETFs). Wealthfront requires an initial investment of $500 [7] and Betterment $10 which creates a much lower entry point for new investors [8].

These robo-advisors are using passive investing strategies to create a streamlined and effective investing approach. Users receive asset allocation based on their risk tolerance and time frame, and then these robo-advisors perform regular rebalancing to maintain those initial allocations. Because of that, robo-advisors do not trade based on short-term factors of an asset (e.g., news, market sentiment, etc.). Investors who are looking to achieve returns exceeding that of the market through active management cannot do so using these robo-advisors.

2.3 AI Stock Recommendations Platforms

Fintech applications which recently became available use AI technology. AI provides recommendations regarding the purchase or stocks based on machine learning algorithms and sophisticated models that analyze stock price movements occurring in the market. These companies such as VectorVest, Helwa, and Simmer do not offer traditional portfolio management as some might expect. Instead, they rely on their recommendation engines which provide investors with AI-generated stock purchase recommendations. VectorVest is one of the examples where a recommendation will be given to the user once there is an opportunity to purchase a stock (e.g. "10 Shares of TSLA"), and the user must go to the brokerage account via the app to place an order and execute the trade. The issues associated with this type of system are as follows:

1) Delay in executing trades: by the time the investor receives the recommendation to execute a trade, market conditions could change.
2) Emotional impact: users who are impacted by emotional issues while making investment decisions will often delay and question recommendations or change the investment at a moment due to uncertainty in the market.
3) No portfolio tracking: most of their alert systems provide users with investment recommendations, but there is no way for the user to track how many of those recommendations have been acted on and have been working within their portfolio or how well or poorly the overall portfolio has performed.

3. METHOLOGY

3.1 The Process

Nexus Gate Fund operates as a five-stage automated trading cycle that executes continuously during market hours. The system is designed to minimize user intervention while maintaining transparency, auditability, and real-time responsiveness to market conditions. Figure 1 illustrates the overall process flow, from data acquisition through trade execution and performance evaluation.

The trading cycle begins with real-time data acquisition from external market sources. Price data for 30 tracked securities is retrieved using the Finnhub API, covering a diverse set of assets including major market indices (SPY, QQQ, DIA, IWM), large-cap equities (AAPL, MSFT, NVDA, TSLA), commodities (GLD, SLV), and leveraged exchange-traded funds (TQQQ, SQQQ, SOXL, SOXS). To comply with API rate limits and ensure system stability, fetched price data is cached locally for 90 seconds before subsequent requests are issued. In parallel, recent market news headlines associated with each tracked security are collected from the same API, including headline text, summaries, sources, and timestamps, providing contextual information for sentiment-driven decision making.

Following data acquisition, the system performs both technical and sentiment-based analysis to identify potential trading opportunities. Technical indicators are computed using historical price data obtained from Yahoo Finance. Specifically, the system calculates the Relative Strength Index (RSI) to detect overbought or oversold conditions, short-term and medium-term simple moving averages (20-day and 50-day SMAs) to evaluate trend direction, and the Moving Average Convergence Divergence (MACD) to assess momentum changes. These indicators collectively generate preliminary buy or sell signals for each security.

Concurrently, sentiment analysis is applied to recent news headlines using a local TextBlob-based classifier. Each headline is converted into a polarity score ranging from −1 (negative) to +1 (positive), and scores are aggregated on a per-ticker basis to determine overall sentiment strength. To reduce noise and limit the volume of data sent to the language model, the system filters securities to identify a subset of actionable tickers. A ticker is considered actionable if it exhibits strong technical signals, corresponds to an existing portfolio position, or demonstrates significant sentiment impact, defined as an absolute polarity score exceeding 0.3. This filtering step substantially reduces token usage while preserving decision-relevant information.

For actionable tickers, the system constructs a compressed prompt containing the current portfolio state, active holdings with unrealized profit or loss, technical indicator outputs, and summarized news sentiment. This prompt is submitted to the GPT-4o model using a low-temperature configuration (0.3) to promote consistent and low-variance decision outputs. The language model returns a structured trading recommendation consisting of a proposed action (BUY, SELL, or HOLD), a target security when applicable, a concise rationale referencing specific market signals or news events, and a confidence score expressed as a percentage. To mitigate risk, trade execution is restricted to recommendations with confidence levels exceeding 70 percent.

When a trade is approved, execution is handled by a rule-based engine that enforces portfolio safety constraints. Buy orders allocate up to 75 percent of available cash and execute purchases at the current market price using whole shares. Sell orders liquidate 50 percent of an existing position, allowing the system to reduce exposure while maintaining partial holdings. Additional safeguards prevent the system from simultaneously holding inverse leveraged ETF pairs and block trades when market price data is invalid or unavailable. After execution, the portfolio state is immediately updated to reflect changes in cash balance, position quantities, average cost basis, and total portfolio value.

All trading decisions, regardless of whether execution occurs, are logged to Google Sheets for auditability and historical analysis. Each log entry records the decision timestamp, technical and sentiment signals for all tracked securities, the AI-generated recommendation and rationale, confidence score, executed share quantity, and updated portfolio metrics. Performance statistics are computed from the complete trading history, including total return, win rate, best and worst trading days, volatility, maximum drawdown, and Sharpe ratio. These metrics are presented through the web dashboard, enabling users to monitor performance without direct system interaction. The full trading cycle repeats every 60 seconds during standard market hours, ensuring timely response to market changes while maintaining compliance with external API constraints.

3.2 Strategy Layer

The strategy layer that we have now revolves around the concept of swing trading. Moreover, if we are doing well, we will be a bit above the Standard and Poor's 500 (S&P 500). On the contrary, if we are doing poorly, we will be a bit below the S&P 500. Essentially, we are following the market but mainly focusing on buying high liquid stocks like SPY, QQQ, etc. Now the way we do this is to ask GPT-40 to do this for us. For instance, we ask AI to buy liquid ETFS when it trades, and this makes sense because at the end of the day, AI is the one making the decisions, and we are simply just giving it data. This is the strategy that we use right now. However, in the next section, we will see why this strategy is not so great.

3.3 LLM Limitations

While large language models (LLMs) such as GPT-4o provide flexible reasoning and natural language generation capabilities, their use in automated trading systems introduces several technical and operational constraints that must be explicitly managed. Nexus Gate Fund incorporates the LLM as a bounded decision-support component rather than an autonomous trading authority, and the system architecture reflects limitations related to numerical reasoning, reliability, latency, and external dependency.

A fundamental limitation of LLMs is their inability to perform symbolic or numerical computation. Rather than executing arithmetic operations, LLMs generate statistically plausible text based on learned patterns, which can lead to errors in basic arithmetic, percentage calculations, and statistical reasoning. In a financial context, this may manifest as incorrect portfolio allocations, misinterpreted magnitudes, or flawed comparisons between numerical values. To mitigate this risk, Nexus Gate Fund delegates all numerical computation to deterministic Python modules. Technical indicators, sentiment scores, profit and loss values, and portfolio metrics are precomputed before being supplied to the LLM, whose role is restricted to qualitative interpretation rather than calculation. Despite this separation, the model may still misinterpret numerical relationships, particularly when values are presented in similar linguistic contexts.

LLMs are also susceptible to hallucinations, producing confident but factually incorrect statements. In automated trading, hallucinations may appear as fabricated causal explanations; misattributed news events, or references to technical patterns do not present in the input data. Such errors are particularly challenging to detect because hallucinated statements are often expressed with high confidence. Nexus Gate Fund mitigates hallucination risk through conservative execution thresholds, objective pre-filtering of actionable securities, and comprehensive decision logging, which enables post-hoc auditing. However, hallucinations remain an inherent limitation of current LLM technology.

Operational constraints further limit system performance. Token-based pricing creates a direct relationship between data volume and cost, making it impractical to transmit full market and news data for all tracked securities at each decision cycle. To address this, the system employs aggressive prompt compression, reducing token usage by approximately 60–80 percent through security filtering, headline truncation, and sentiment aggregation. While effective in controlling cost and latency, this approach may exclude potentially relevant but weakly signaled information.

Latency introduced by external API calls also constrains system behavior. LLM inference typically incurs delays of several seconds, during which asset prices may change significantly in volatile markets. As a result, Nexus Gate Fund operates on a seconds-to-minutes timescale, in contrast to traditional quantitative hedge funds that execute trades at microsecond latencies. Consequently, the system is designed for swing trading strategies rather than high-frequency trading.

Additional limitations arise from the non-deterministic nature of LLM outputs. Even with low-temperature settings, identical inputs may yield different decisions across executions, complicating back testing and performance attribution. Moreover, the internal reasoning process remains opaque, limiting interpretability and systematic strategy refinement. The model's fixed training cutoff further restricts adaptability, as it lacks direct awareness of post-training market developments and cannot learn from its own trading history without fine-tuning.

Finally, system robustness is affected by structured output fragility and external API dependency. Deviations from expected response formats may lead to parsing failures, in which case the system defaults to conservative HOLD behavior. Reliance on third-party APIs also introduces exposure to outages, rate limits, pricing changes, or model deprecation, all of which can disrupt trading operations.

Collectively, these limitations constrain Nexus Gate Fund to medium-frequency trading of highly liquid securities, where modest execution delays are acceptable and where the LLM's ability to synthesize technical indicators and news sentiment provides value. Future work may explore hybrid architectures that combine deterministic algorithmic trading for routine decisions with LLM-driven reasoning for complex, news-driven scenarios.

3.4 Future implementations

Most quants use statistical modeling to drive their strategy layer. Moreover, they assume that the market follows Arithmetic Brownian Motion and they base their statistical models based on the null hypothesis that in:

$$r_{t+1} = \propto + \beta \phi_t + \varepsilon_{t+1}$$

$$H0: \beta = 0$$

$$H1: \beta \neq 0$$

This means that in a market where returns are iid, β is 0 which means that yesterday tells us nothing about today and that we are just left with:

$$r_{t+1} = \propto\ + \varepsilon_{t+1}$$

which gives no edge because we are just left with expected return (r_{t+1}) = drift ($\propto$) + volatility (ε_{t+1}).

On the contrary, if β is not 0, then that means yesterday does tell us something about today, and that we can predict where the market will go based on what we have today (ϕ_t) and its effect on the market (β). This is what many quants aim to do. Moreover, they aim to reject the null hypothesis H0 and side with the alternative hypothesis H1. This is assuming that we are using a Neyman-Pearson test.

Now, when comparing Nexus to this model, these quant models are deterministic. Moreover, ϕ_t includes facts such as z-score and sentiment. These values are not calculated by AI but rather there are functions that do this math.

Next semester, the goal is to involve the option to buy high liquid ETFs, and other functions. This will allow us to spot deviations in β and hopefully achieve a high return (r_{t+1}).

Now, aside from the strategy layer, we also aim to have a fully decoupled system and a multi-user platform by the end of the next semester. Not only will we have a decoupled system and a multiuser system, but we will also connect to Alpaca, which is a broker, and this allows users to make real trade using real money.

4. RESULTS

In this section, we will present and evaluate the results of our automated trading bot. We have collected trading data from the bot from early July to the current day. This model currently trades with paper money. The model had a starting portfolio value of $10,000.

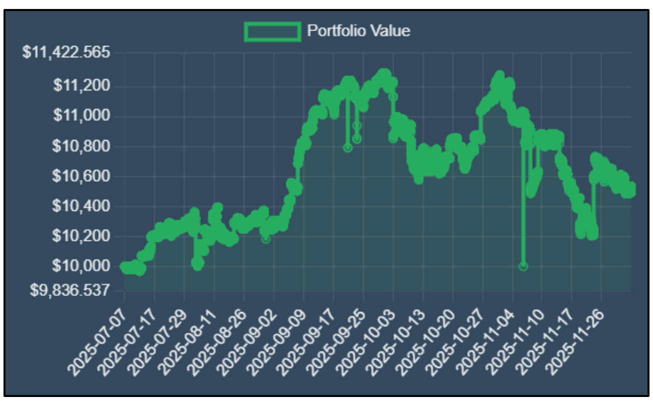

The graph shows an upwards trend at the beginning and peaked at around $11,200. An error in our algorithm occurred between November 4th and November 10th, causing the portfolio to be reset. We were able to fix this error and allowed the trading bot to resume its original portfolio value before the value was reset. Although the trading bot has not reached its peak again, we will continue to refine and work on the algorithm for the trading bot to increase portfolio value.

The picture above shows the data that belongs to our main graph. The most important aspects of the image above are volatility and Sharpe ratio. Volatility describes how rapid the changes are. According to the data above, the volatility is high, which is expected since we are swing trading. For the Sharpe ratio, we have -1.26 and this basically tells us that even though we achieved 10,536.067, it was a very bumpy ride getting there. Now, this is expected for this graph as this graph is an intraday representation of our model. Moreover, the noise of the market affects our intraday trade. So, it is not guaranteed that we will continuously make profit but rather, it is expected that we will experience significant drops because the noise of the market affects our return.

The intraday representation of the graph is not a good indicator of telling us how well our bot is doing. To fix this issue, we use the end of the day snapshot graph which is below:

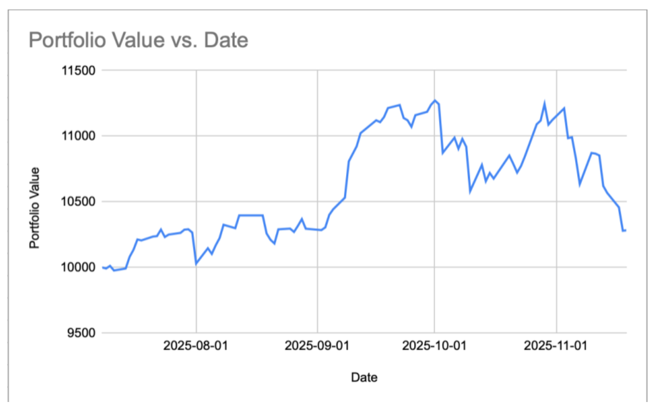

The graph above has a Sharpe ratio of 0.55 and a volatility of 1.60. When looking at the Sharpe ratio of this graph, it is an improvement compared to the intraday graph. This makes sense because there is almost no noise at the end of the day. Also, the volatility is lower, which shows that the bot is not making rapid changes. From this graph we cannot say that we are beating S&P 500. This is due to how the S&P 500's Sharpe is 0.45 while our Sharpe is 0.55. However, our sample size is not big enough to make this claim, and we need to continue to run this to get a larger sample size.

The graph above used to be a lot bigger. Furthermore, we used to have a large sample size until a few days ago when we lost all data from Nov 14th to Dec 10th. This was caused by a malfunction in the end-of-the-day snapshot function, and this just goes to show the errors in the system. However, with our future implementations being decoupled architecture, these errors will not happen anymore, and we will be able to accumulate a much larger time series of Nexus' returns.

5. CONCLUSION

This paper presents Nexus Gate Fund, an automated trading bot that handles trading stocks on behalf of the user to generate revenue. Compared to traditional robo-advisors, Nexus Gate Fund is more accessible to university students. The proposed algorithm illustrates the logic behind the decisions made by artificial intelligence on when to buy, hold, or sell a stock. We will continue to refine the algorithm to increase revenue that the automated bot can make.

6. ACKNOWLEDGMENTS

Our thanks to ACM SIGCHI for allowing us to modify the templates they have developed. We would also like to thank the California State University, Fullerton Engineering and Computer Science department, their faculty advisors, and Dr. Yu Bai.

7. REFERENCES

[1] Yunduan Lou, Pu Sun, Yifeng Yu, Shangping Ren, and Yu Bai. 2025. TT-DSC: Enhancing YOLO for Marine Ecosystem through Efficient Tensor Train-based Depthwise Separable Deep Neural Network. ACM Trans. Auton. Adapt. Syst. Just Accepted (May 2025). https://doi.org/10.1145/3735138

[2] X. Ma, P. Sun, S. Luo, Q. Peng, R. F. DeMara and Y. Bai, "Binarized l1 -Regularization Parameters Enhanced Stripe-Wise Optimization Algorithm for Efficient Neural Network Optimization," in IEEE Journal of Emerging and Selected Topics in Industrial Electronics, vol. 5, no. 2, pp. 790-799, April 2024, doi: 10.1109/JESTIE.2023.3313050.

[3] QuantifiedStrategies. 2025. Decoding the Medallion Fund: What We Know About Its Annual Returns. Retrieved from https://www.quantifiedstrategies.com/decoding-the-medallion-fund-what-we-know-about-its-annual-returns/

[4] Wikipedia. 2025. Renaissance Technologies. Retrieved from https://en.wikipedia.org/wiki/Renaissance_Technologies

[5] SmartAsset. 2025. What Are Typical Investment Minimums in Hedge Funds? Retrieved from https://smartasset.com/investing/hedge-fund-minimum-investment

[6] Arielle O'Shea and Dayana Yochim. 2025. Betterment vs. Wealthfront: 2025 Comparison. NerdWallet. Retrieved from https://www.nerdwallet.com/investing/learn/betterment-vs-wealthfront

[7] FinanceBuzz. 2025. Wealthfront vs. Betterment: Two Ways to Grow Your Money. Retrieved from https://financebuzz.com/wealthfront-vs-betterment

[8] Davila, H. and Hilliard, J. (2024) Financial Literacy and Student Debt: Survey of College Students, Social Science Research Network. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4836687 (Accessed: 16 December 2025).