Computational AI Model for Finance Application: A Cognitive Stock Analysis Architecture using Large Language Model

Nha Van Thanh Do
Georgia Institute of Technology
ndo46@gatech.edu

***Abstract*—This project investigates how a Large Language Model can model bounded rational decision-making by integrating sentiment-driven heuristics with multi-step reasoning to generate reliable stock market recommendations under resource and information constraints. Drawing on foundational cognitive science theories: bounded rationality, dual-process theory, mental representations and cognitive heuristics. This tool was designed and implemented that mirrors key aspects of human financial cognition. The system processes real-time market data and news through a dual-process pipeline: a fast, heuristic-based sentiment analysis module (System 1) and a deliberate, multi-step reasoning engine (System 2), both powered by a locally deployed Gemma3:12b model using Ollama open source model. The tool enforces bounded rationality through a hard 10-article limit per stock and constrained token budgets, reflecting human cognitive limitations. Results show that the dual-process architecture produces more robust and balanced recommendations than either system alone, and that explicit information constraints improve decision clarity rather than degrade it. These findings contribute to the understanding of how cognitive principles can be computationally operationalized, offering implications for both cognitive science theory and practical AI-assisted decision support.**

I. INTRODUCTION

A. *Research Question*

How can an AI model bounded rational decision-making by integrating sentiment-driven heuristics with multi-step reasoning to generate concise and reliable recommendations under resource and information constraints?

B. *Importance and Significance*

This research question addresses a fundamental challenge in cognitive science: understanding how minds make decisions under real-world constraints. Herbert Simon introduced the concept of bounded rationality, arguing that human decision-makers operate with limited cognitive resources and, instead of optimizing, individuals satisfice seeking solution that meet minimum aspiration levels rather than maximum potential [1]. This is especially evident in the stock market, where it is virtually impossible to process all available information, buy at the exact dip, and sell at the peak. Investors must balance multiple factors and accept trade-offs to reach decisions that are at least profitable.

Recently, Binz and Schulz demonstrated that LLMs exhibit human-like cognitive biases, suggesting they may serve as useful models of human reasoning [2]. Their work showed that GPT-class models replicate well-known cognitive phenomena such as anchoring and framing effects. This finding opens a compelling research topic: if LLMs reason in ways structurally similar to human cognition, then building a system that deliberately leverages these cognitive parallels can both advance our theoretical understanding and produce a practical decision-support tool.

Financial decision-making is a particularly rich domain for cognitive science investigation because it inherently involves uncertainty, incomplete information, time pressure, and affective responses—all factors central to theories of human cognition. Kahneman's dual-process framework distinguishes between fast, intuitive judgments (System 1) and slow, deliberate analysis (System 2), both of which are engaged when investors evaluate market conditions [3]. Furthermore, Newell and Simon's work on mental representations demonstrates that effective problem-solving requires converting perceptual input into structured symbolic forms that can be manipulated [4]. A stock analyst, for example, does not reason over raw news text; instead, they mentally encode news as bullish or bearish signals, risk factors, and catalysts—exactly the kind of structured representation this system produces.

C. *Relation to Cognitive Science*

This project is grounded in for central cognitive science concepts:

- Bounded relationality: The system enforces hard constraints on information processing with 10 articles per stock, limited token budgets, and simplified decision categories (Buy/Hold/Sell) mirroring how humans satisfice rather than optimize under cognitive limitations.
- Mental representation: The system transforms unstructured perceptual input (raw news articles) into structured symbolic representations (sentiment scores, risk factors, catalysts) that enable higher-order reasoning.
- Dual-process theory: The design includes two parallel processing systems: System 1 for fast sentiment-based heuristics providing immediate affective responses to

news) and system 2 for deliberate multi-step reasoning analyzing fundaments, risks, and catalysts.

- Cognitive heuristics: The system implements recency weighting (prioritizing recent news), price anchoring (using current price as a reference for target price estimation) and availability heuristics (relying on readily accessible information).

II. MODEL/TOOL DESIGN

A. Model Build and Tech Stack

This is a AI-powered data analysis tool built on top of Gemma3:12B Ollama open-source model and some free APIs for news and articles retrieval from yfinance for market data, Google News/ Yahoo Finance RSS feeds for news articles. There are some limitations with free API service because I cannot get the more accurate and thoughtful analysis from experts like Wall Street Journal or New York Times. But since this tool is for proof of concept purpose, Google News and Yahoo Finance are good enough and produce reliable actions. It also depends on how I prompt and provide instruction to LLM to optimize and enhance the accuracy. A good prompt could help model to perform solid analysis under constraints. Because of that, Gemma3: 12B models can also provide sufficient reasoning capability while remaining computationally tractable on current hardware. There were some attempts to use more powerful model Llama 3.3-70B using free Groq API, however, I reached the token limitation when fetching 10 news articles and multiple stocks at the same time.

The project is Python-based modular architecture with Streamlit UI for user interaction. It's running locally to avoid cloud dependencies and token limitations.

Component	**Tech Stack**	**Purpose**
LLM	Gemma3: 12B	Local LLM
Backup LLM	Llama-3.3-70B	Cloud-based LLM
Market Data	Yfinance API	Real-time stock prices
News Aggregation	Google News, Yahoo Finance	Stock-specific, general market news
RSS Parsing	Feedparser + BeautifulSoup	HTML/RSS content extraction
Frontend	Streamlit	Interactive web-based application
Programming Language	Python	

B. Cognitive Science Concepts in System Design

Concept	Design Element	Detail
Bounded relationality	10 article limit per stock	ARTICLE_COUNT = 10 500-1000 tokens per LLM call
Mental representations	Structured intermediate data	Raw news -> JSON with sentiment_score from -10 to 10, catalysts, risk_factors
Dual-process theory	2 parallel systems to evaluate stock's news	System1: sentiment analysis (temp=0.2) System 2: recommendation generation (temp=0.3)
Cognitive Heuristics	Recency weighting	Recent news prioritized. Current price anchors target price
	Recommendation	Buy/Hold/Sell

C. Source of Principles

- Bounded rationality from Simon's work on how real decision-makers operate under constraints of limited information, memory, and time [1]
- Mental representations from Newell and Simon's theory that human cognition converts perceptual input into symbolic representations for manipulation and reasoning [4].
- Dual-process theory from Kahneman distinguishing between fast, intuitive (System 1) and slow, deliberate (System 2) cognitive processes [3].
- Cognitive heuristics from Tversky and Kahneman's identification of mental shortcuts (anchoring, availability, recency) that humans use for complex judgments under uncertainty [5].
- The insight that LLMs can serve as cognitive models comes from Binz and Schulz's empirical demonstration that GPT-class models exhibit human-like biases and reasoning patterns [2]. Even though, I'm not using GPT models but Gemma and Llama can be considered as comparable models to GPT.

D. System Implementation

The system follows a three-step hierarchical pipeline that mirrors the flow of human cognitive processing:

Step 1: Data Collection and Perception

- The system accepts stock symbols from the user and gathers market information using a multi-method fallback approach for robustness: Market indices (S&P 500, NASDAQ, VIX) via yfinance for macro context; stock fundamentals (price, PE ratio, sector, volume) via yfinance; news artciles (up to 10 per stock) through yfinance and Google News RSS feed.
- The fallback design ensures data availability even when individual sources fail-analogous to how humans rely on multiple information channels.

Step 2: System 1 – Fast Heuristic Processing (Sentiment Analysis)

- For each stock, this function (or system) aggregates news headlines and summaries (truncated to 200 characters per article); sends them to LLM with a low temperature setting (0.2) for deterministic, focused classification; produces a structured JSON output containing sentiment score (-10 to 10), sentiment label (Bullist/Neutral/Bearish), key points, catalysts, and risk factors.
- The low temperature setting reflects the fast, automatic nature of System 1 thinking – rapid pattern recognition without extensive deliberation. The system prompt constraints the LLM to act as "a financial analyst specializing in news sentiment analysis", narrowing its response space and minimize hallucination.

Step 3: System 2 – Deliberate Reasoning (Market Overview and Recommendations)

- Market-level overview: The system synthesizes market indices and general news into a coherent market narrative (temperature = 0.3, allowing slightly more creative synthesis). This mirrors how a human analyst first forms a macro view before evaluating individual positions.
- Per-Stock Recommendation Generation: This will combine System 1 sentiment outputs with fundamental data to produce detailed recommendations including: action (Strong Buy/ Moderate Buy/ Hold/ Sell), target price, rationale (5-8 sentences), risk level, timeframe, and news catalyst. This step explicitly integrates the fast heuristic output (sentiment) with deliberate analysis modeling the interplay between System 1 and System 2.

Fallback heuristics:

- This is applied when the LLM fails to produce a valid recommendation. The fallback multiplier will be used to infer the target price based on sentiment direction. This mirrors how humans rely on simpler mental shortcuts when more complex reasoning fails or is unavailable.
- Bullish sentiment: 8% upside multiplier, target = price * 1.08
- Bearish sentiment: 5% downside multiplier, target = price * 0.95
- Neutral sentiment: 2% upside multiplier, target = price * 1.02

After the entire process is completed, recommendations are sorted by potential return in descending order, reflecting salience bias – humans naturally attend to the most attention highest opportunity first. The system generates portfolio health assessments, alerts for bearish positions, and actionable items. Output is displayed through a Streamlit web dashboard featuring expandable per-stock sections, color-coded sentiment indicators, and styled recommendation cards.

E. System Architecture Diagram

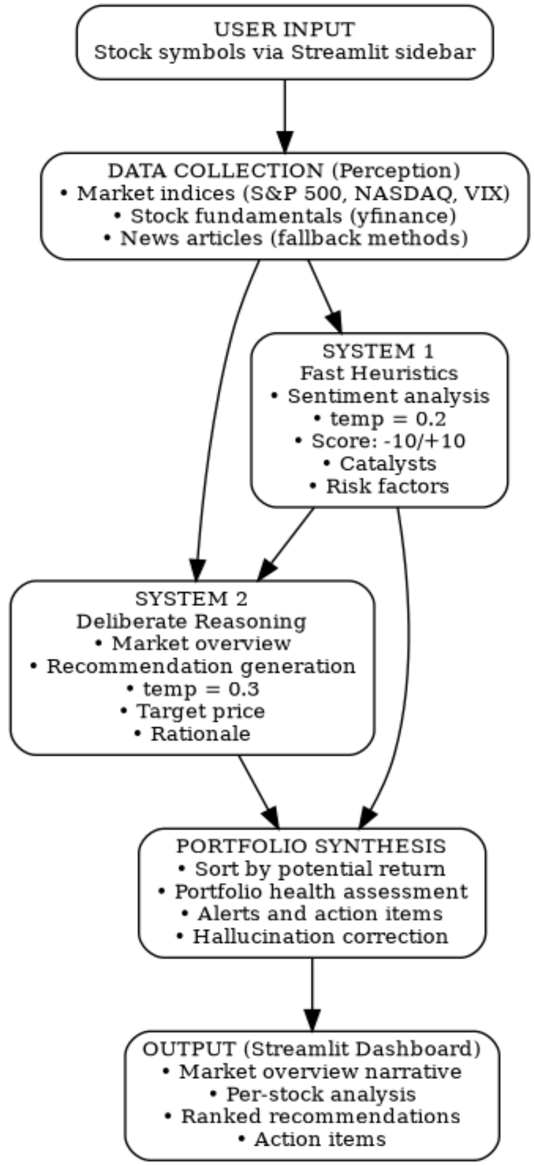

Fig. 1. System Architecture Diagram. Generated by ChatGPT

III. RESULTS

A. Findings from Building and Evaluating the Tool

- Bounded Rationality Constraints Improve Decision Clarity: One of the most counterintuitive results was that the 10 articles limit per stock did not degrade decision quality. During development, I was experimenting with varying article counts: 3,5,10,15. While 3 and 5 could be fed to Groq API key, it's just too low especially during the volatility time. 15 is just too high and lower the LLM performance because of extra noise and extra compute resources. At 10 articles, the system consistently produced focused, defensible recommendations. This directly supports Simon's theory that satisficing within bounds can be as effective as unconstrained optimization

[1]. Cognitive constraints force prioritization of relevant information and what matters the most.

- Dual Process Architecture produces more robust decisions: The parallel operation of System 1 (fast sentiment heuristics) and System 2 (deliberate reasoning) consistently produced more balanced recommendations than either system alone. System 1 alone tended to be reactive—overly bullish on positive news and overly bearish on negative news without considering fundamentals. System 2 alone, without the sentiment signal, sometimes produced recommendations disconnected from recent market events. When both systems contributed, the final recommendation reflected both the immediate affective response to news (System 1) and the reasoned assessment of fundamentals and risk (System 2). The combination of two parallel systems is important because they compensate for each other and harden the final decisions, supporting Kahneman's dual-process framework [3].

- Mental Representations are essential for Multi-step reasoning: A critical design discovery was that the LLM could not reliably perform multi-step reasoning over raw news text. When raw articles were directly fed into the recommendation prompt, the model frequently produced incoherent or hallucinated outputs. However, when news was first transformed into structured representations like sentiment scores, categorized catalysts, and enumerated risk factors, the recommendation quality improved substantially. This validates Newell and Simon's theory that symbolic representations are necessary preconditions for effective problem-solving [4]. Raw data cannot guide reasoning; intermediate symbolic representations are necessary for coherent decision-making. The transformation pipeline (raw articles to structured JSON to reasoned decisions) directly implements the mental representation framework.

- Heuristic Fallbacks demonstrate cognitive resilience: The anchoring-based fallback mechanism (8% bullish, 5% bearish, 2% neutral multipliers) activated when the LLM failed to produce valid JSON output. Notably, these simple heuristic recommendations were often directionally correct; they captured the general market sentiment even without detailed analysis. This mirrors how human decision-makers fall back on simpler heuristics when cognitive resources are exhausted or when complexity exceeds processing capacity. Recency weighting is particularly effective since the latest news always reflects the most recent activity and more accurate information, enabling faster processing with acceptable accuracy.

- Temperate as a proxy for cognitive effort: The use of different LLM temperature settings for System 1 (0.2) and System 2 (0.3) proved to be an effective proxy for cognitive effort. Lower temperature produced more deterministic, focused outputs appropriate for quick classification tasks (sentiment analysis), while slightly higher temperature allowed more creative synthesis appropriate for generating market narratives and nuanced recommendations. This mapping of temperature to cognitive processing mode offers a novel computational operationalization of the System 1/System 2 distinction.

- Prompt engineering as cognitive instruction: Using the more powerful Llama 70B, the model could understand context better and provide accurate analysis with just a simple prompt. However, there is token limitation made this infeasible to analyze 10 articles and multiple stocks. This forced me to use smaller model downloaded locally and fit with my laptop configuration. With Gemma 12B, it requires to completely redesign the prompts to be more specific to avoid hallucination. This parallels how human experts need less explicit instruction than novices. The smaller model, with less "knowledge," requires more structured guidance to perform effectively. A well-crafted prompt helps the model perform solid analysis under constraints, much like how clear mental frameworks guide human reasoning under uncertainty.

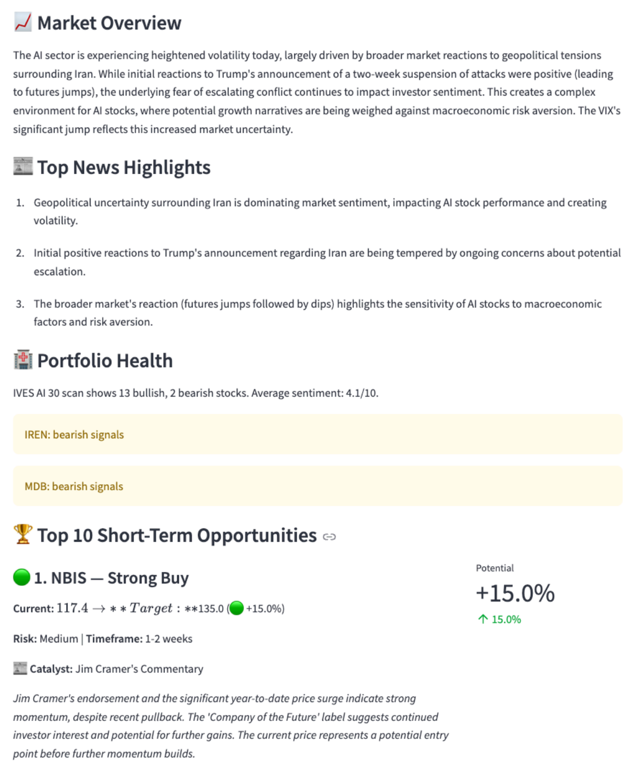

Fig. 2. Market Overview and Highest Returned Stock

B. System Behavior and Performance Patterns

The was tested across a portfolio of 12 representative stocks spanning multiple sectors: SBUX, MSFT, SNOW, NVDA, TSLA, UNH, JEQP, NFLX, T, AAPL, META, GOOGL and also from the list of 30 AI companies in IVES index. The system will analyze all stocks and return the top 10 promising stocks along with the analysis and recommendation.

Key behavioral observations:

- Sentiment-price correlation: The sentiment extraction module successfully correlates with market sentiment indicators, validating the mental representation approach. Stocks with strongly bullish sentiment scores (>5) were typically associated with positive recent price movements, and vice versa for bearish scores.

- Processing time: Single stock analysis takes around 40-60s while IVES analysis on 30 AI stocks usually takes 15-20 mins on local hardware.
- Hallucination management: Despite prompt engineering, the LLM occasionally generates unsupported claims about stock performance. The price hallucination detection mechanism (flagging when LLM-generated prices deviated >50% from actual prices) and output validation with schema enforcement helped mitigate but did not fully eliminate this issue.
- Graceful degradation: When news was unavailable for certain stocks (particularly smaller-cap names), the system defaulted to neutral sentiment with appropriate uncertainty markers rather than producing unreliable analysis.

IV. DISCUSSION

A. *Implicatioins for Cognitive Science*

- Bounded rationality is computationally tractable. Simon's theory is often discussed in abstract terms, but this project shows it can be implemented as concrete engineering constraints (article limits, token budgets, simplified categories) that measurably affect system behavior. The finding that 10 articles outperformed 20 in decision clarity provides computational evidence for the satisficing hypothesis—more information does not necessarily lead to better decisions [1].
- Dual-process theory has architectural implications. The success of the System 1/System 2 architecture suggests that separating fast heuristic processing from slow deliberate reasoning is not merely a descriptive framework but a prescriptive design principle. AI systems that implement both processing modes produce outputs that are both responsive to immediate context (via System 1) and grounded in careful analysis (via System 2). Both systems operating in parallel and compensating for each other produce qualitatively better outputs than either independently [3].
- Mental representations are a prerequisite for reasoning. The failure of the LLM to reason over raw text and its success when given structured representations provides computational support for Newell and Simon's mental representation hypothesis [4]. This finding has implications beyond financial analysis: any AI system that needs to perform multi-step reasoning may benefit from an explicit representation formation stage that mirrors human cognitive encoding.
- Practical constraints drive innovation. Token limitations forced a pivot from cloud-based Llama 3.3-70B to local Gemma3:12b, which unexpectedly improved system reliability and eliminated external dependencies. This mirrors how human cognitive constraints can lead to adaptive, often superior solutions, a phenomenon consistent with the ecological rationality perspective that constraints and environmental structure jointly shape effective decision strategies.

B. *Why This System Matters*

First, it bridges cognitive science theory and practical AI engineering. By grounding each system component in a specific cognitive science principle, the project demonstrates that theoretical frameworks from cognitive psychology can guide the design of AI systems in principled ways—not as metaphors, but as engineering specifications.

Second, it contributes to the emerging field of studying LLMs as cognitive models. Following Binz and Schulz's demonstration that LLMs exhibit human-like biases [2], this project shows that LLMs can be deliberately structured to implement specific cognitive architectures, opening avenues for using AI systems as testbeds for cognitive theories.

Third, it addresses a real-world need. As a full-time employee and part-time graduate student, my time is limited and it becomes increasingly difficult to follow the market, especially during periods of stock market uncertainty. This tool is designed to replace manual human effort in gathering and processing financial information, performing thoughtful analysis and multi-step reasoning similar to a human analyst. The system demonstrates that cognitively-inspired design can produce useful, actionable outputs for investors who face the same bounded rationality constraints that the system models.

C. *Real-World Applications*

The cognitive architecture developed in this project could be applied to several real-world contexts:

- Personal financial advisory: The tool directly serves as an AI assistant for individual investors, providing structured analysis and recommendations that respect human cognitive limitations and deliver true value by streamlining manual processes.
- Decision support in high uncertainty domains: The dual process architecture could be adapted for medical diagnosis, legal analysis, or intelligence assessment, any domain where both fast pattern recognition and deliberate reasoning are valuable.
- Educational tool for cognitive science: The system could serve as a pedagogical demonstration of how bounded rationality, dual process theory, and mental representations function in practice.
- AI system design: The cognitive principles used here (information constraints, dual processing, structured representations) could inform the design of more human aligned AI systems across domains. The insight that prompt engineering quality matters more for smaller models mirrors how structured guidance matters more for less experienced human reasoners.

V. CONCLUSION

A. What I Learned

This project provided deep insight into the relationship between cognitive science and AI system design. The most important lesson was that cognitive constraints are not obstacles to be overcome but they are design principles that can improve system performance. The 10 articles limit, simplified decision categories, and structured representations all began as approximations of human cognitive limitations, but each proved to be a valuable architectural choice that improved the quality and clarity of the system's outputs.

I also learned that dual process theory translates remarkably well into a computational architecture. The separation of fast sentiment analysis (System 1) and deliberate recommendation generation (System 2) produced consistently better results than either approach alone, providing practical validation of Kahneman's theoretical framework.

On the technical side, the project deepened my understanding of LLM deployment, prompt engineering, and the importance of output validation. The pivot from cloud APIs (Llama 3.3-70B via Groq) to a local model (Gemma3:12b) was driven by practical token constraints but ultimately improved system reliability and control. This experience demonstrated that practical constraints can drive innovation—a lesson that itself reflects the cognitive science principle that constraints shape adaptive behavior.

Furthermore, I learned the importance of the relationship between AI and cognitive science—how AI operates under cognitive concepts and how to apply these theories to further enhance and improve AI agent applications. The project also helped me sharpen my technical skills and build a meaningful personal AI assistant product to streamline manual processes and improve daily productivity.

B. Key Insights

- Bounded rationality is a feature, explicit information constraints improve decision clarity.
- Dual process architectures produce more robust decisions than single-process approaches.
- Mental representations are essential for multi-step LLM reasoning.
- Cognitive heuristics provide reliable fallback mechanisms when full analysis fails.
- LLM temperature can serve as a computational proxy for cognitive effort.
- Prompt engineering quality matters more for smaller models, paralleling how structured guidance matters more for novice human reasoners.

C. Future Directions

- Continuous testing engine: Compare system recommendations against actual market outcomes over time to measure predictive accuracy, moving beyond decision robustness to financial performance validation
- Feedback loops: Implement learning mechanisms that adjust heuristic parameters based on past recommendation performance, modeling human experiential learning
- Multi-agent architecture: Extend the system with specialized agents such as risk agent, momentum agent to model the diversity of expert perspectives in group decision-making
- Adaptive bounded rationality: Dynamically adjust the article limit based on market volatility, processing more information during uncertain periods and less during stable ones, modeling how humans allocate attention adaptively
- Premium data sources: Integrate expert analysis from sources like Wall Street Journal or New York Times for richer and more accurate mental representations

VI. LIMITATIONS

A. System Limitations

- Limited validation: The system was not fully evaluated against actual future stock performance. Because the recommendation is for 1-4-week timeframe and due to global situation tension recently, stock markets have experienced significant volatility that lowers the model accuracy. In addition, original goal was to predict actual stock prices, but this proved too ambitious and requires much more resources and market understanding than just 10 articles. The evaluation methodology was pivoted to focus on recommendation consistency, sentiment-signal correlation, and decision robustness rather than predictive accuracy.
- LLM hallucination: Despite prompt engineering, the LLM occasionally generates unsupported claims about stock performance. Mitigation through output validation and schema enforcement helped but did not eliminate this completely. Using the smaller Gemma 12B requires more specific prompts to avoid hallucination compared to more powerful models.
- News availability bias: Inconsistent news availability for smaller stocks limits analysis scope; some stocks yield fewer than 10 relevant articles. The system depends on freely available news sources (Yahoo Finance, Google News RSS), which cannot match the quality and depth of premium sources like Wall Street Journal or New York Times.
- Single-model dependency: The system relies on Gemma3:12b for all cognitive processing. While sufficient for proof of concept, a more robust implementation would use ensemble methods or multiple models with different strengths.
- No formal user study: The system's impact on human decision-making was not empirically tested with human participants. Claims about cognitive support are based on

system behavior analysis rather than user experience data.

- Static heuristic parameters: The fallback multipliers (8%, 5%, 2%) are fixed rather than learned or adapted, limiting the system's ability to improve over time.
- Computational constraints: While feasible on standard hardware, full local inference is slower than API calls (~40–60 seconds per stock), making real-time analysis of large portfolios impractical.

B. *Future Improvement*

- Implement a historical checking module to validate recommendations against historical market data.
- Add ensemble methods combining multiple LLMs for more robust reasoning.
- Develop adaptive heuristic parameters that learn from past recommendation outcomes.
- Integrate premium data sources (SEC filings, earnings transcripts, expert analysis) for richer mental representations.
- Implement real-time streaming analysis for time-sensitive market events.
- Add explainability features that trace each recommendation back to the specific cognitive process (System 1 or System 2) that generated it.
- Explore more powerful local models as hardware improves, reducing the need for extensive prompt engineering.

References

[1] S. Herbert, "A Behavioral Model of Rational Choice", The Quarterly Journal of Economics, vol. 69, no.1, pp. 99-118, 1955.

[2] B. Marcel, E. Schulz, "Using Cognitive Psychology to Understand GPT-3", Proceedings of the National Academy of Sciences, vol. 120, no. 6,2023

[3] D. Kahneman, "Thinking, Fast and Slow", New York, NY, USA, Farrar, Straus and Giroux, 2011

[4] N. Allen, S. Herbert, " Human Problem Solving". Englewood Cliffs, NJ: Prentice-Hall, 1972

[5] A. Tversky, D. Kahneman, "Judgment under uncertainty: Heuristics and biases", Science, vol. 185, no. 4157, pp.1124-1131, 1974.

VII. Appendix

A. *RESEARCH PLAN SUMMARY*

Week	Task	Description	Hours	Complete
3	1	Research and decide project track	3	Y
3	2	Understand cognitive science concepts used in the project	3	Y
4	3	Draft version of each concept and sketch a plan to apply in project	3	Y
4	4	Outline system architecture plan and LLM models	2	Y
4	5	Prepare, write, organize and complete project pitch	5	Y
Total Hours			16	
Project Pitch Due				
5	6	Implement sentiment module	3	Y
5	7	Implement System 1	4	Y
5	8	Implement System 2	4	Y
6	9	Bounded rationality parameter controls	2	Y
6	10	Hallucination monitoring module	1	Y
6	11	Integration testing of pipelines	2	Y
6	12	Define evaluation metrics	1	Y
6	13	Create sentiment test scenarios	1	Y
7	14	Build automated evaluation pipeline	4	Y
7	15	Design comparison experiments	2	Y
7	16	Run baseline experiments	2	Y
8	17	Run bounded rationality experiments	2	Y
8	18	Compare system 1 and system 2 outputs	2	Y
8	19	Initial Midpoint Check-in Report	3	Y
8	20	Final Check-in Midpoint Report	3	Y
Total Hours			36	
Optional Midpoint Check-in Due				

9	21	Analyze experimental data	4	Y
9	22	Statistical comparison of results	4	Y
9	23	Interpret cognitive implications	4	Y
9	24	Draft analysis notes	3	Y
10	25	Write methods section	3	Y
10	26	Write evaluation section	3	Y
10	27	Finalize technical details	4	Y
11	28	Final data validation	3	Y
11	29	Enhance UI	6	Y
11	30	Final bug fix or improvement	5	Y
12	31	Gathering data and start initial final report	5	Y
12	32	Finalize report	4	Y
13	33	Final validation for both application and report	5	Y
13	34	Submit final report materials	1	Y
Total Hours			54	
Final Report Due				
Total Hours for the entire project			106	
Final Presentation Due				

B. Github Repository

https://github.gatech.edu/ndo46/CS6795_SP26_TERM_PROJECT_FINANCIAL_ANALYST