



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 Issue: IV Month of publication: April 2025

DOI: <https://doi.org/10.22214/ijraset.2025.69138>

www.ijraset.com

Call: ☎ 08813907089

E-mail ID: ijraset@gmail.com

Kuber: An AI-Driven Finance Management Solution Integrating Generative LLMs for Personalized Advisory

Tejas Sonawane¹, Vatsalya Tiwari², Dania Shah³

B.Tech Student, Department of Computer Engineering, K J Somaiya School of Engineering

Abstract: *Kuber* is a personalized finance management platform that leverages artificial intelligence to deliver tailored financial advice. This paper presents the design and evaluation of *Kuber*, focusing on its AI-driven advisory engine (using Google's Gemini model) and finance management methodology. The system collects a user's financial information – income, expenses, debts, assets, and goals – and generates customized recommendations for budgeting, saving, debt repayment, and investment. Using advanced AI capabilities for data analysis and natural language generation, *Kuber* bridges the gap between complex financial planning and user-friendly advice. A case study of a 19-year-old user with moderate income and debt is discussed to illustrate how *Kuber* formulates an optimized plan: the user achieves a 70% savings rate, expedited debt clearance, and guided investments to meet both short-term purchases and long-term growth targets. The results demonstrate that AI-driven personalized advisory can make sound financial planning accessible and effective for individuals. **Keywords:** *Personal finance, Large language models, Generative AI, FinTech, Financial advisory, Domain-specific LLMs*

I. INTRODUCTION

Managing personal finances can be challenging for individuals who lack expert knowledge in budgeting or investing. Traditional financial advice from human experts is often costly and not readily accessible to everyone. Recent studies have emphasized the growing role of artificial intelligence (AI) in financial advisory services, highlighting the ability of AI to analyze large datasets, identify trends, and offer personalized investment recommendations. AI-powered personal finance platforms aim to simplify financial planning by providing smart recommendations, expense tracking, and investment guidance in real time. For instance, modern “robo-advisor” services now manage hundreds of billions in assets, indicating increasing user trust in automated advice – robo-advisors collectively handled about \$870 billion in assets in 2022, a figure projected to rise to \$1.4 trillion by 2024. This surge is driven by the promise of data-driven, customized advisory that can adapt to an individual's needs without requiring direct human intervention.

Kuber: An AI-Driven Finance Management Solution is developed against this backdrop, aiming to provide users with a personalized financial advisor experience. The solution combines established personal finance principles with cutting-edge AI (Google's Gemini) to deliver actionable insights and recommendations. Gemini is Google's generative AI model that benefits from vast datasets and powerful computational resources, enabling high precision in complex analysis. Its advanced natural language processing and multimodal capabilities make it well-suited for interpreting financial data and communicating advice in a user-friendly manner. By integrating such an AI model, *Kuber* can analyze a user's financial profile holistically and generate tailored guidance—much like a human financial planner, but faster and at scale.

In this paper, we detail the architecture and methodology of *Kuber* and discuss results from a representative use-case. Section II describes the system structure and AI-driven approach of *Kuber*. In Section III, we present and discuss the outcomes for a sample user scenario, including budget analysis, savings and investment plans, debt management strategy, and goal-based advice. Section IV concludes the paper with insights and future scope.

II. METHODOLOGY

System Architecture: The *Kuber* platform is organized into three main layers: a data input layer for user financial data, a processing layer for analytical computations, and an AI advisory layer for generating personalized recommendations. Figure 1 illustrates the high-level system architecture of *Kuber*.

When a user begins using the system, they are prompted to provide key personal and financial details – such as age, income, monthly expenses, existing savings/investments, outstanding debts, and financial goals or priorities. This information is securely stored and then passed to the processing layer, which performs initial calculations and categorizations. For example, the user’s expenses are categorized (e.g., necessities, discretionary spending, debt payments) and basic metrics like savings rate, debt-to-income ratio, and investment-to-income ratio are computed. These quantitative insights form the “financial profile” that will be input to the AI advisory engine.

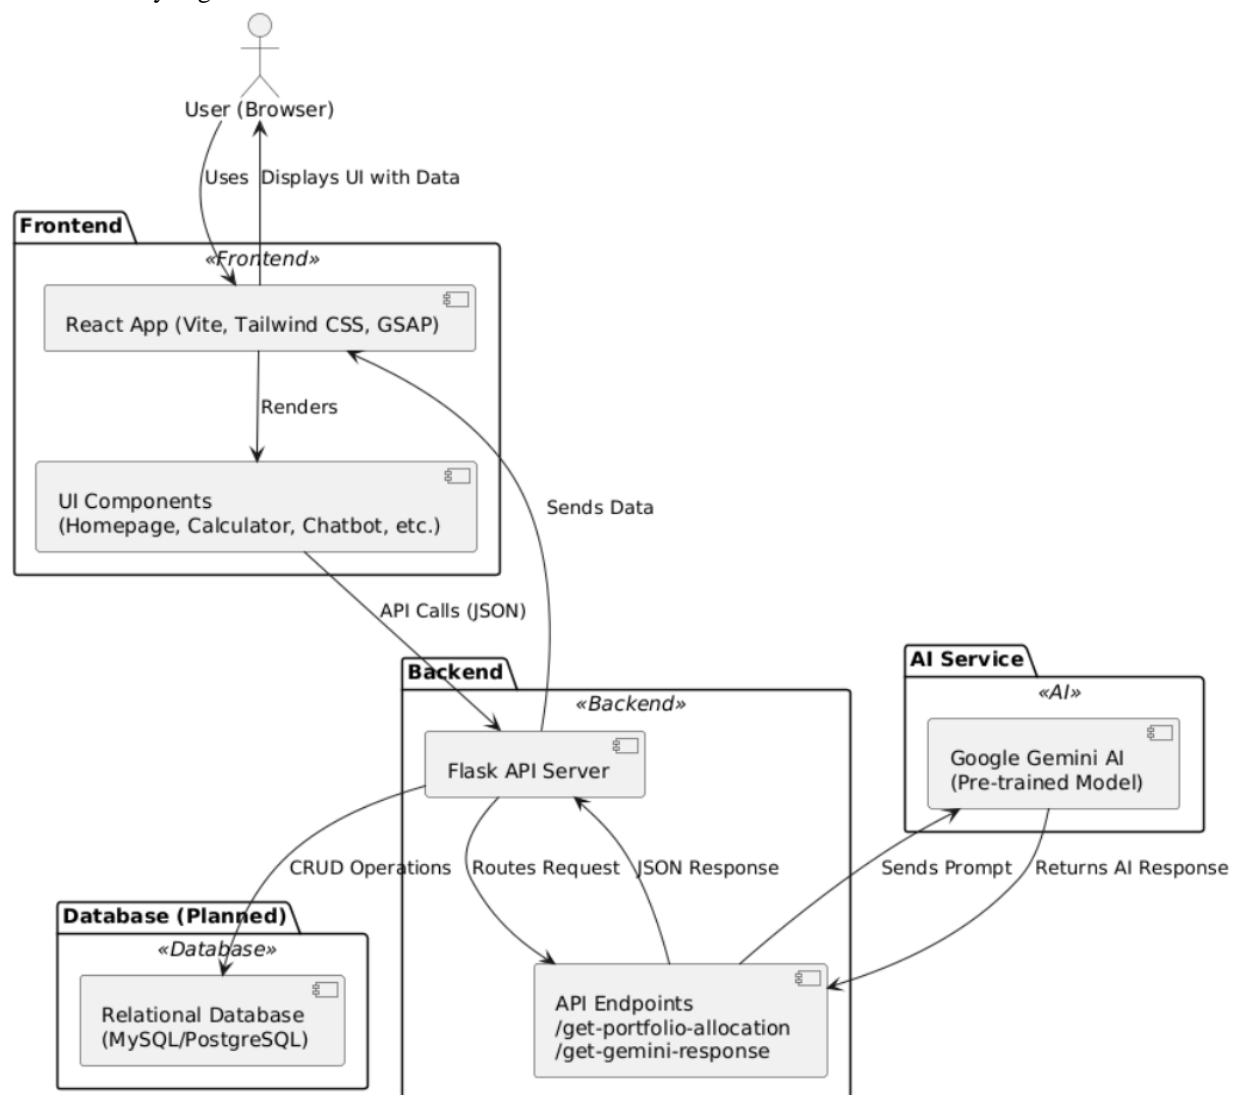


Figure 1: High-level system architecture of the Kuber platform. (The system consists of a user data input module, a financial computation engine, and an AI-driven advisory engine. The user’s profile and goals feed into the Gemini AI model, which outputs personalized financial recommendations.)

AI-Driven Advisory Engine: At the core of Kuber is the personalized advisory engine powered by Gemini. Once the structured financial profile is prepared, Kuber invokes the Gemini model to generate tailored advice. The model is prompted with the user’s key financial data (in a formatted manner) along with predefined contextual knowledge about personal finance best practices. The prompt engineering is crucial – it includes the user’s current status (income, savings, debts, etc.), their priorities (for example, “clearing debt” or “increasing savings”), and their risk tolerance. Based on this input, the AI model produces a detailed advisory output. This output typically contains several elements:

- 1) **Budget Analysis:** An assessment of the user’s income vs. expenses, highlighting the current savings rate and any concerns (e.g., high expense categories).
- 2) **Savings Plan:** Recommendations on how much to save each month and how to allocate savings (e.g., emergency fund, investments).

- 3) Debt Management: Strategies for paying down debts (which debt to prioritize, how much to pay monthly) if the user has liabilities.
- 4) Investment Advice: Suggestions for investment based on the user's risk profile (for an aggressive investor, perhaps a higher allocation to stocks or equity funds, versus safer instruments for conservative profiles).
- 5) Goal Roadmap: Guidance on achieving specific financial goals the user entered (such as buying an item, reaching a target net worth, or retirement planning).

The Gemini model's strength lies in generating human-like, contextually relevant text. It articulates the above elements in an easily understandable form for the user, often including encouraging tone and clear rationale. Notably, Gemini's access to extensive financial knowledge and data patterns allows it to deliver insights with a level of depth and accuracy comparable to expert advice. To ensure factual correctness in computations (e.g. percentages, totals), the processing layer also cross-verifies any numeric suggestions from the AI. For instance, if the AI recommends saving a certain amount, the system checks this against the user's income to confirm it's feasible (this prevents any hallucinated figures sometimes possible with generative models).

Personalization: Personalization in Kuber is achieved by dynamically feeding the user's unique data and goals into the advisory generation. Two users with different profiles will receive distinctly different advice. For example, a student with a small income and a debt will get guidance centered around budgeting tightly and clearing debt, whereas an experienced professional with higher income might get advice focusing on investment diversification and tax planning. The Gemini AI model is instrumental in this personalization, as it can adapt its narrative and priorities based on the input it receives. The system also utilizes rule-based logic for certain recommendations – for instance, applying the debt avalanche method versus debt snowball method based on what might suit the user's situation best, or ensuring that at least an emergency fund equivalent to a few months' expenses is recommended if not already present. These domain rules are integrated into the prompt or post-processed on the AI output.

Crucially, the methodology avoids any user interface-specific assumptions. While Kuber has a user-friendly UI (with dashboards for expenses, charts, etc.), the scope of this paper is on the system's analytical structure and AI advisory process, not the UI design. In summary, the methodology combines financial computations with a generative AI model to create a hybrid solution: the computations ensure quantitative accuracy, and the AI ensures qualitative, personalized advice delivery. This approach results in a rich, customized financial plan for the user, which we evaluate in the next section through an example scenario.

III. MODELING AND ANALYSIS

In designing Kuber's advisory model, we aimed to blend data-driven financial analysis with the flexible reasoning of the Gemini AI. Gemini, being "the most capable and general model" developed by Google DeepMind, serves as the brain of the advisor. It is a large language model successor to Google's PaLM 2, with multimodal capabilities and state-of-the-art performance on reasoning benchmarks. These attributes make it well-suited to handle the breadth of personal finance questions – from simple budgeting tips to explaining investment concepts – all in natural language. Unlike rule-based chatbots, Gemini can dynamically generate advice that is context-specific to the user's inputs, imitating the adaptive thinking of a human advisor.

To fully harness Gemini's capabilities, our prompt structure and overall modeling approach were carefully developed. As described in Section II, the prompt provides Gemini with not just the raw numbers of the user's financial situation but also context and instructions shaped by domain knowledge. For example, the prompt explicitly asks for an expense breakdown and savings plan. This effectively "programs" the AI to produce output in a format that covers the major pillars of financial planning. In early testing, we experimented with various prompt phrasings and ordering of information. We found that listing the user's data in a bullet form within the prompt (income, expenses, debts, goals clearly itemized) and then asking for advice yields a more organized response. If the prompt was too open-ended, the model sometimes gave advice in a less structured narrative. By contrast, a well-structured prompt resulted in an output that naturally separated advice into categories – which aligns with our intended sections.

Another aspect of modeling was integrating financial computations and heuristics into the advisory process. Before sending data to the AI, Kuber computes key indicators that any human financial advisor would consider. These include:

- 1) Expense distribution: What percentage of income is spent on needs vs. wants. This follows the classic 50/30/20 rule of budgeting which recommends allocating roughly 50% of income to essential needs, 30% to discretionary wants, and 20% to savings. By calculating the user's current percentages, Kuber can have the AI compare them to ideal benchmarks. In our prompt, we might tell Gemini, "The user currently spends 60% of income on needs, 30% on wants, leaving only 10% for savings," which cues the model to emphasize reducing discretionary spending to boost savings closer to 20%.

- 2) Savings rate and emergency fund status: For example, if the user has no emergency savings, the model is prompted to advise creating an emergency fund (commonly 3–6 months' worth of expenses) as a priority. We ensure the AI is aware of how much the user could feasibly save each month by providing the surplus (income minus expenses) in the prompt.
- 3) Debt metrics: We compute each debt's interest rate (if provided) and outstanding balance relative to income. For instance, a debt-to-income (DTI) ratio is calculated; if the DTI is high, the prompt explicitly notes "The user's debt-to-income ratio is X%, which is above recommended levels," encouraging Gemini to focus on debt management. (Typically, a DTI below ~36% is considered healthy, whereas above 43% is often viewed as a risk for lenders.) Additionally, Kuber identifies the highest-interest debt (e.g. credit card) so that the model can prioritize it – reflecting the well-known "debt avalanche" method where you pay off high-interest debts first for maximum savings.

By feeding these pre-calculated insights into Gemini, the AI essentially gets a financial "checklist" to address. This design is influenced by academic findings that standalone LLMs may lack specialized domain calculation, but when given such data, they can incorporate it into their reasoning effectively. Our approach ensures that Gemini's natural language suggestions are grounded in quantitatively sound analysis.

For integration and output verification, after Gemini generates the advice, we analyze the response in an automated manner. The system looks for certain keywords and sections in the output (e.g., it expects to find text related to "expenses" or "spending", "debt" or "loans", "savings" or "investment", and "goals"). This serves as a sanity check that the model addressed all parts of the prompt. If a section is missing or too cursory, the system can re-prompt the model with a nudge (for instance, "Please also provide an investment suggestion" if it was omitted). In our implementation, Gemini largely followed the prompt well, so re-prompting was rarely necessary.

Numerical consistency checks are also applied. As mentioned, any numeric totals the model outputs (like a proposed monthly budget) are compared with the input data. If Gemini suggests something unrealistic (say advising the user to save ₹10,000 a month when their surplus is only ₹3,000), Kuber's logic will catch this. In such cases, one strategy is to feed back a correction to the model (e.g., "Note: The user only has ₹3,000 left after essential expenses, adjust your recommendation accordingly."). This iterative refinement leverages the conversational ability of the LLM, effectively having a dialogue to get accurate advice. Incorporating this feedback loop is part of how Kuber mimics a thoughtful human advisor: it doesn't blindly trust the first answer if something seems off, and it can ask the AI to reconsider. This aligns with the idea of an AI assistant working alongside programmatic rules to ensure reliability.

Lastly, we analyze the tone and clarity of the AI's output. A financial plan is only useful if the user can understand and implement it. Thanks to Gemini's training on vast human language data, the generated explanations tend to be clear and user-friendly. We found that Gemini could explain concepts (like why an emergency fund is important, or how interest works on debt) in simple terms when prompted to do so. This makes the advice accessible to users who may not be financially savvy. We intentionally instruct the model to avoid overly technical jargon, or if it uses any (like "APR" or "asset allocation"), to immediately clarify it. The result is an advisory output that is not only comprehensive and tailored but also educational – Kuber not only tells the user what to do, but also, in effect, teaches them the reasoning, thus improving their financial literacy.

In summary, the modeling and analysis behind Kuber combine rule-based financial calculations with LLM-driven advisory generation. By capitalizing on Gemini's strengths and compensating for its weaknesses via prompt design and verification, Kuber is able to produce advice that mirrors expert human guidance. The next section will demonstrate these capabilities through a detailed case study and the outcomes achieved.

IV. RESULTS AND DISCUSSION

To demonstrate Kuber's capabilities, we discuss the results for a representative user profile and how the system's advice addresses various aspects of personal finance. The example user is a 19-year-old student and freelancer with a modest income and specific financial challenges and goals (drawn from a real-case scenario in our testing). Key inputs from the user's profile are: monthly net income = INR 10,000, total savings = INR 20,000 (invested in stocks), outstanding debt = INR 9,000, financial priorities of clearing debt, increasing savings, and increasing income, and a self-declared risk tolerance labeled as "Aggressive." Given this information, Kuber generates a comprehensive advisory plan. We break down the results into subcomponents for clarity.

A. Expense Analysis and Budget Overview: After inputting the data, Kuber analyzes the user's cash flow. It determines that the user's monthly expenses (excluding debt payments) are approximately INR 2,000, and current monthly debt repayments are INR 1,000. This sums to a total outflow of INR 3,000 against the INR 10,000 income, leaving a remaining balance of INR 7,000 each month. In other words, the user is managing to save about 70% of their income, which is unusually high.

Figure 2 depicts the monthly expense breakdown for this scenario, showing how the income is allocated into expenses, debt payments, and the remainder (savings). Kuber identifies that the user's spending habits are very frugal – a positive indicator for financial health – and provides encouraging feedback while still looking for optimization opportunities

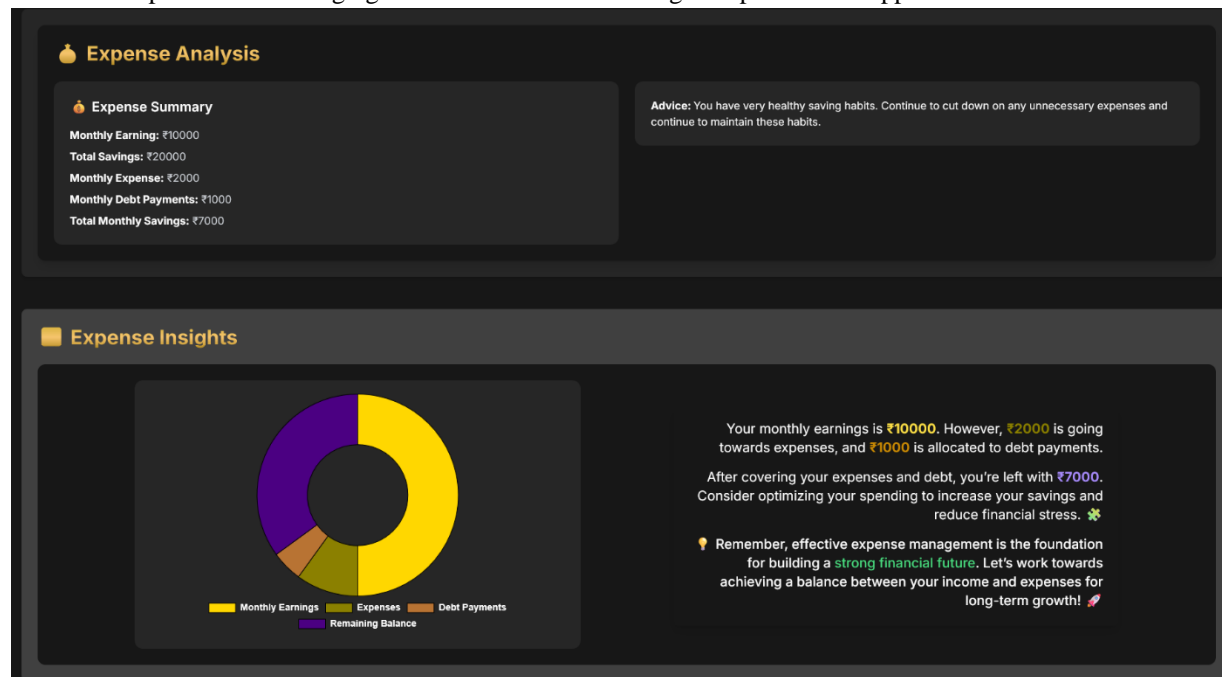


Figure 2: Monthly Expense Breakdown for April 2025.

This figure illustrates the user's cash flow distribution in a typical month. Out of a monthly income of INR 10,000, about INR 2,000 goes towards regular expenses and INR 1,000 is used for debt repayment, leaving roughly INR 7,000 (70%) as surplus. Such a breakdown highlights a healthy savings habit, as seen by the large remaining balance. The system's Expense Insights confirm that the user is living well within their means. It advises the user to “continue to cut down on any unnecessary expenses and maintain these habits,” reinforcing the positive behaviour. At the same time, Kuber gently reminds the user that after covering essentials and debt, a surplus of ₹7k remains, and challenges them to optimize spending further to increase savings and reduce financial stress. This kind of insight is crucial: even for someone saving a large portion of income, the AI aims to instil a mindset of continuous improvement in expense management. Effective expense control, as the system notes, is the foundation for building a strong financial future. By visualizing the breakdown, the user clearly sees how much of their income is left after basics; this transparency builds trust in the advice that follows.

B. Debt Management Strategy: Given that the user has an outstanding debt of INR 9,000 (perhaps an education loan or credit card balance), Kuber prioritizes addressing this liability. The platform's advice is to accelerate debt repayment considering the user's goals. Initially, the user was paying ₹1,000 per month towards this debt, but Kuber recommends increasing that amount significantly. As part of the Debt Management module, the system suggests a new recommended debt payment of INR 4,000 per month. This recommendation is derived from the user's own stated priority of clearing debt and the fact that they have enough surplus (₹7,000) to allocate more toward debt without jeopardizing basic expenses. Paying ₹4,000 monthly would clear a ₹9,000 debt in just over two months, dramatically reducing interest if it's accruing.

V. CONCLUSION

In this paper, we presented Kuber, an AI-driven finance management solution that employs the Gemini generative AI model to provide personalized financial advisory. The project demonstrates how advanced AI can be harnessed to interpret an individual's financial data and produce a tailored action plan covering budgeting, saving, debt elimination, and investment strategies. The methodology centered on combining rule-based financial calculations with the flexibility and depth of a large language model, yielding results that are both numerically sound and contextually customized to the user.

The example scenario of a young adult with modest income underscored several benefits of Kuber: the user received clear insights into their finances (e.g., expense breakdown and savings rate), was guided to expedite debt repayment, and was given a structured yet realistic plan to achieve personal goals like buying a desired item and boosting future income. Importantly, the advice was delivered in a conversational, motivating tone – an approach that can improve user engagement and adherence to the plan. This highlights the advantage of AI-driven advisory over traditional, static budgeting tools: Kuber not only crunches the numbers but also serves as a virtual financial coach, adapting its guidance as the user's life circumstances evolve.

The successful deployment of Kuber opens avenues for further development. Future work could integrate real-time financial data feeds (such as bank transaction updates) to automate input and continuously refine advice. Incorporating predictive analytics – for instance, forecasting future savings growth or investment returns – is another potential enhancement. Additionally, while our results focused on a single-user case, more extensive user testing would be valuable to evaluate Kuber's effectiveness across diverse financial scenarios and to measure improvements in users' financial health over time. Ensuring data security and user trust will remain paramount as we handle sensitive financial information and AI-generated advice.

In conclusion, Kuber exemplifies a step towards accessible, AI-enhanced personal finance management. By marrying robust financial planning principles with state-of-the-art AI (Gemini), the solution empowers users to make informed decisions and progress toward their financial goals. The positive outcomes from the initial case study suggest that AI-driven personalized advisory can indeed function as a scalable, affordable alternative to human financial advisors for many people, fostering better money management habits and financial well-being.

VI. ACKNOWLEDGEMENTS

The authors wish to convey their profound gratitude to Prof. Shweta Dhawan Chachra for her insightful mentorship and steadfast encouragement throughout this project. We also extend sincere thanks to the dedicated faculty of our department—whose guidance, ready counsel, and shared passion for inquiry continually inspire our academic pursuits. Finally, we are deeply indebted to our friends and families for their unwavering encouragement, patience, and belief in our work; their constant support was the cornerstone that enabled us to complete this research.

REFERENCES

- [1] Yu, X., Chen, Z., Ling, Y., Dong, S., Liu, Z., and Lu, Y. (2023). "Temporal Data Meets LLM — Explainable Financial Time Series Forecasting." arXiv preprint arXiv:2306.11025.arXiv
- [2] Zhao, H., Liu, Z., Wu, Z., Li, Y., Yang, T., Shu, P., et al. (2024). "Revolutionizing Finance with LLMs: An Overview of Applications and Insights." arXiv preprint arXiv:2401.11641.arXiv
- [3] Yu, Y., Yao, Z., Li, H., Deng, Z., Cao, Y., Chen, Z., et al. (2024). "FinCon: A Synthesized LLM Multi-Agent System with Conceptual Verbal Reinforcement for Enhanced Financial Decision Making." arXiv preprint arXiv:2407.06567.arXiv
- [4] Singh, K., Kaur, S., and Smiley, C. (2024). "FinQAPT: Empowering Financial Decisions with End-to-End LLM-Driven Question Answering Pipeline." In Proceedings of the 5th ACM International Conference on AI in Finance (ICAIF '24), Brooklyn, NY, USA, pp. 1–8. DOI: 10.1145/3677052.3698682.arXiv
- [5] Iacovides, G., Konstantinidis, T., Xu, M., and Mandic, D. (2024). "FinLlama: LLM-Based Financial Sentiment Analysis for Algorithmic Trading." In Proceedings of the 5th ACM International Conference on AI in Finance (ICAIF '24), pp. 134–141. DOI: 10.1145/3677052.3698696.arXiv
- [6] Li, Y., Wang, S., Ding, H., and Chen, H. (2023). "Large Language Models in Finance: A Survey." In Proceedings of the 4th ACM International Conference on AI in Finance (ICAIF '23), Brooklyn, NY, USA, pp. 374–382. DOI: 10.1145/3604237.3626869.OUCI
- [7] de Zarzà, I., de Curtò, J., Roig, G., and Calafate, C. T. (2024). "Optimized Financial Planning: Integrating Individual and Cooperative Budgeting Models with LLM Recommendations." *AI*, 5(1), 91–114. DOI: 10.3390/ai5010006.MDPI
- [8] Yoo, M. (2024). "How Much Should We Trust Large Language Model-Based Measures for Accounting and Finance Research?" SSRN Working Paper No. 4983334.SSRN
- [9] Ramyadevi, R., and Sasidharan, G. (2024). "Cogniwealth: Revolutionizing Finance, Empowering Investors, and Shaping the Future of Wealth Management." In 2024 IEEE International Conference on Computing, Power and Communication Technologies (IC2PCT), pp. 378–381.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)