FinPersona: An LLM-Driven Conversational Agent for Personalized Financial Advising

Takehiro Takayanagi[1,2], Masahiro Suzuki[1], Kiyoshi Izumi[1], Javier Sanz-Cruzado[2], Richard McCreadie[2], and Iadh Ounis[2]

[1] The University of Tokyo, Tokyo, Japan
takayanagi-takehiro590@g.ecc.u-tokyo.ac.jp
izumi@sys.t.u-tokyo.ac.jp
[2] University of Glasgow, Glasgow, Scotland, United Kingdom
{javier.sanz-cruzadopuig,richard.mccreadie,iadh.ounis}@glasgow.ac.uk

Abstract. Conversational financial assistance is gaining increasing attention as a subdomain within the broader field of conversational AI, attracting interest from both the AI and finance communities. The application of conversational agents in finances is a particularly challenging subfield to explore, as financial scenarios require personalization and emotional support, especially during periods of market turmoil. In this paper, we present **FinPersona**, a Large Language Model (LLM)-driven conversational agent for personalized financial advising. FinPersona enables a non-expert user to explore investment options through conversation and dialogue driven by knowledge of the individual user and their circumstances. Through FinPersona, we aim to illustrate the salient features of a working prototype, as well as raise awareness of the many challenges unique to the financial domain. A video demonstration of the system is available at: https://youtube/tvJAYEg73Ro

Keywords: Financial Advisory · Conversational Agent · Personalization

1 Introduction

Personalized advice plays an invaluable role in our daily lives, particularly when it comes to making decisions in critical areas such as finance and healthcare. Advisors in these fields offer personalized guidance that considers the advisee's specific preferences and circumstances. In addition, they often offer emotional support, helping clients navigate the stress and uncertainty that can accompany these decisions. However, these advisory services often come at a high cost, making them unaffordable for a large sector of the population. To mitigate this issue, automated financial decision support systems have been widely studied, with a primary emphasis on investment-related predictions, such as predicting asset profitability and investors' preferences [13,15,16,17].

Recent studies suggest that the main added-value of the financial advisors lies in their ability to offer personalization and emotional support rather than convey traditional financial indicators such as investment returns [2,6][3]. Personalization

[3] For instance, a survey from Dimensional's 2017 Global Investor Feedback, conducted on 18,967 individual investors, shows that investors value a sense of security, peace

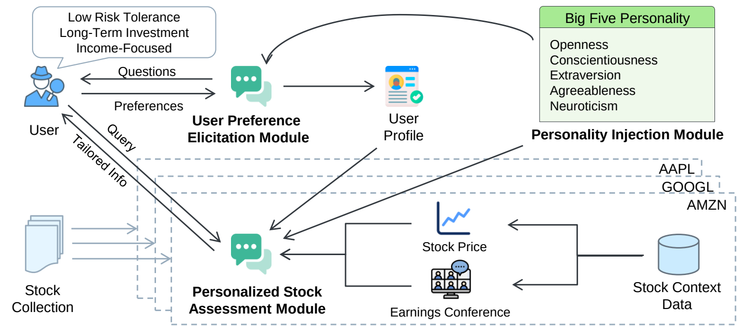

Fig. 1. FinPersona Architecture

involves adapting the provided financial advice to the client's current situation (current wealth, employment), needs and preferences (risk aversion, sector preferences), while emotional support helps clients manage their emotions and prevent them from making rushed decisions, particularly during periods of market volatility. Indeed, if LLMs are to be leveraged in this domain then they need the capability to provide effective emotional support in addition to (accurately) answering the client's questions [10]. This makes LLMs for finance a particularly rich and difficult research area, since LLMs not only need to process numerical and temporal information (which they are often poor at), but they also need to manage the user's state of mind, as well as adapt the way they convey information to the background and expertise of the user. For this reason, there has been a surge in interest from both the AI and finance communities into LLM-based financial advising [8,10]. However, to our knowledge, no system has been developed that uses LLMs to provide both personalized advice and emotional support to investors. Therefore, in this paper, we introduce **FinPersona**, a new financial advisor platform powered by LLMs, which can support investors' financial decision-making. Our platform is able to integrate preferences information from investors to enable personalization. It also provides emotional support capabilities to enable the conversational agent to adopt personalities that align with the needs of investors, as suggested by [10]. The primary audience for the FinPersona platform are PhD students and researchers working on information retrieval (IR) and recommendation, who might be interested in the finance space as a new domain to explore.

2 The FinPersona Platform

FinPersona is a conversational assistant capable of dynamically collecting and capturing the users' preferences, as well as analyzing and integrating contextual information to provide useful and personalized insights to investors. As illustrated in Figure 1, FinPersona consists of three modules: first, users can begin by choosing the personality they want to assign to their personal chatbot

of mind, and knowledge of their personal financial situation more than investment returns when interacting with financial advisor services.

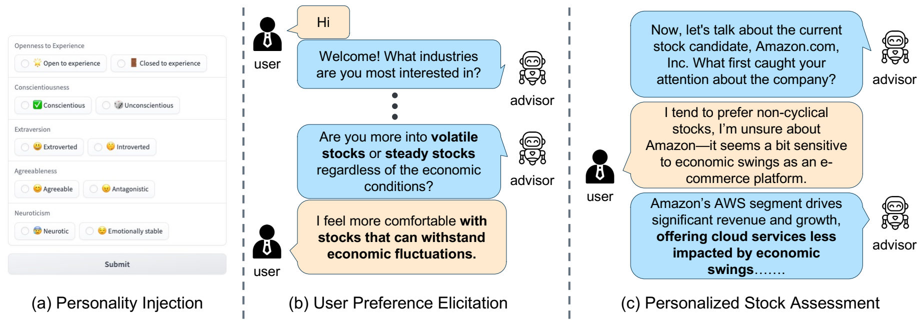

Fig. 2. Conceptual Overview of the FinPersona Support System

through the personality injection module. Next, the advisor dynamically collects the users' preferences through conversations with the user in the user preference elicitation module, thereby building a user profile. Finally, users interact with the advisor to decide whether to invest in a given stock, with the advisor providing tailored information in the personalized stock assessment module. The system has been designed to support large-scale user studies, and has been already successfully tested in preliminary experiments with a group of 20 participants.

Personality Injection Module: The personality injection module aims to simulate a specific personality in an LLM-based financial advisor. To provide support for the users in choosing an advisor acting in a manner aligned with their expectations, we allow a user to select the personality they expect from a financial advisor, based on the Big Five personality traits [3]: openness to experience, conscientiousness, extroversion, agreeableness and neuroticism. As shown in Figure 2(a), for each trait, we allow the user to select between two opposite values. The personality injection module compiles the different traits and injects the selected personality into the conversational financial agent. Recent studies have proposed various methods to induce personality in LLMs through prompting or fine-tuning [4,5,14]. In our FinPersona platform, we adopted and tailored the prompting method by Jiang *et al.* (2024) for assigning a Big Five personality trait to the LLM agent [3] due to its simplicity and effectiveness.

User Preference Elicitation Module: In order to tailor the financial advice to their clients, financial advisors need to collect relevant information about these clients: e.g., their goals, preferences, current situation, etc. In FinPersona, the user's preferences elicitation module is the component dedicated to such a task. It allows the agent to collect relevant information about the users through an interactive multi-turn conversation as shown in Figure 2(b). The advisor proactively asks the user about their preferences while clarifying any concepts unknown to the investor (i.e., user). FinPersona continues the conversation until the LLM-based financial advisor assesses that it has gathered sufficient information from the user. FinPersona follows a predefined list of salient financial indicators that the advisor requires to understand the needs of their clients in the user preference elicitation module. Following the idea of style investing [1], these financial indicators include the clients' preferences for a number of stock features such as their industry, stock style, consistency on dividend payments, and sensitivity to global market changes. These predefined set of financial indicators can be further expanded in FinPersona to include other important investment traits such as the investor's risk tolerance or their financial outlook plans and horizons.

Personalized Stock Assessment Module: The last component of FinPersona enables users to decide whether a target stock is suitable for them by conversing with a LLM-based financial advisor who provides personalized insights about the stock. In FinPersona, the agent discusses only one target stock at a time (e.g., the user explores whether Amazon is a suitable investment, rather than choosing from multiple stock candidates) We choose to focus the conversation on a single stock since this allows us to evaluate the user's financial decision-making process in a more fine-grained manner. Additionally, incorporating multiple stocks would lead to longer prompts for the LLM, leading to difficulties in its ability to handle them effectively [9]. Figure 2(c) shows an example conversation with a user.

In this module, the advisor uses two types of data: first, the user profile collected during the user's preferences elicitation stage; second, multi-modal contextual information about the target stock, such as the stock prices and the earnings summaries [12][4]. Earnings calls are hosted by publicly traded companies to discuss key aspects of their earning reports and future goals with financial analysts and investors [11]. Our approach in FinPersona can be easily expanded into a retrieval-augmented generation (RAG) [7] setting, where the advisor retrieves relevant data from external resources, which will be of interest to the users.

3 Conclusions

In this paper, we propose an LLM-based financial advisory system, FinPersona, designed to provide both personalized advice and emotional support to investors. The system is not intended to recommend specific stocks or financial instruments but rather to provide personalized information to aid users in making informed decisions. FinPersona has been designed as a research toolkit to explore and evaluate the effectiveness of LLM-powered financial advisors and follows strict ethical guidelines.[5] We have conducted an initial evaluation of FinPersona with 20 external users. Our preliminary results show that LLM personalization enables more accurate decisions from the users, although it does not necessarily lead to a higher satisfaction with the advices given. We plan to further evaluate the effectiveness of personalization and emotional support in the future.

[4] We collect earnings conference call transcripts from Seeking Alpha for the last quarter of 2023 and also gather monthly stock prices from 2023 via Yahoo! Finance. Naturally, this data has to be continuously collected over time as it becomes available to ensure that the agent has access to the latest available information.

[5] **Ethical considerations:** Since the system uses LLMs, there is a risk of generating inaccurate information. As we clearly convey to the users of FinPersona, users should verify all information and consult professional financial advisors before making any investment decisions.

Acknowledgments. This work was supported by JST SPRING, Grant Number JP-MJSP2108.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Barberis, N., Shleifer, A.: Style Investing. Journal of Financial Economics **68**(2), 161–199 (2003)
2. Bennyhoff, D.G., Kinniry Jr, F.M., DiJoseph, M.A.: The Evolution of Vanguard Advisor's Alpha®: From Portfolios to People. Tech. rep., The Vanguard Group, Valley Forge, Pensylvannia, USA (2018)
3. Costa Jr., P.T., McCrae, R.R.: A Five-Factor Theory of Personality. In: Handbook of personality: Theory and research, pp. 139–153. Guildford Press New York, 2nd edn. (1999)
4. Jiang, G., Xu, M., Zhu, S.C., Han, W., Zhang, C., Zhu, Y.: Evaluating and inducing personality in pre-trained language models. In: Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS 2023). New Orleans, Louisiana, USA (2024)
5. Jiang, H., Zhang, X., Cao, X., Breazeal, C., Roy, D., Kabbara, J.: PersonaLLM: Investigating the Ability of Large Language Models to Express Personality Traits. In: Duh, K., Gomez, H., Bethard, S. (eds.) Findings of the Association for Computational Linguistics: NAACL 2024. pp. 3605–3627. Association for Computational Linguistics, Mexico City, Mexico (2024). https://doi.org/10.18653/v1/2024.findings-naacl.229
6. Kinniry Jr., F.M., Jaconetti., C.M., DiJoseph., M.A., Zilbering., Y., Bennyhoff., D.G., Yarwood., G.: Putting a value on your value: Quantifying Vanguard Adviser's Alpha in the UK. Tech. rep., The Vanguard Group, Valley Forge, Pensylvannia, USA (2020)
7. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.t., Rocktäschel, T., Sebastian, R., Kiela, D.: Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In: Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS 2020). pp. 9459–9474. Online event (2020)
8. Li, Y., Wang, S., Ding, H., Chen, H.: Large Language Models in Finance: A Survey. In: Proceedings of the 4th ACM International Conference on AI in Finance (ICAIF 2023). pp. 374–382. New York, New York, USA (2023). https://doi.org/10.1145/3604237.3626869
9. Liu, N.F., Lin, K., Hewitt, J., Paranjape, A., Bevilacqua, M., Petroni, F., Liang, P.: Lost in the middle: How language models use long contexts. Transactions of the Association for Computational Linguistics **12**, 157–173 (2024)
10. Lo, A.W., Ross, J.: Can ChatGPT Plan Your Retirement?: Generative AI and Financial Advice. Harvard Data Science Review (2024). https://doi.org/https://doi.org/10.1162/99608f92.ec74a002
11. Medya, S., Rasoolinejad, M., Yang, Y., Uzzi, B.: An Exploratory Study of Stock Price Movements from Earnings Calls. In: Companion Proceedings of the Web Conference 2022 (WWW 2022). p. 20–31. Association for Computing Machinery, Lyon, France (2022). https://doi.org/10.1145/3487553.3524205

12. Mukherjee, R., Bohra, A., Banerjee, A., Sharma, S., Hegde, M., Shaikh, A., Shrivastava, S., Dasgupta, K., Ganguly, N., Ghosh, S., Goyal, P.: ECTSum: A New Benchmark Dataset For Bullet Point Summarization of Long Earnings Call Transcripts. In: Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP 2022). pp. 10893–10906. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates (2022). https://doi.org/10.18653/v1/2022.emnlp-main.748
13. Sanz-Cruzado, J., Richards, E., McCreadie, R.: FAR-AI: A Modular Platform for Investment Recommendation in the Financial Domain. In: Proceedings of the 46th European Conference on Information Retrieval (ECIR 2024), Part V. p. 267–271. Springer-Verlag, Glasgow, United Kingdom (2024). https://doi.org/10.1007/978-3-031-56069-9_30
14. Shao, Y., Li, L., Dai, J., Qiu, X.: Character-LLM: A Trainable Agent for Role-Playing. In: Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023). pp. 13153–13187. Association for Computational Linguistics, Singapore (2023)
15. Takayanagi, T., Chen, C.C., Izumi, K.: Personalized dynamic recommender system for investors. In: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 2246–2250 (2023)
16. Takayanagi, T., Izumi, K.: Incorporating domain-specific traits into personality-aware recommendations for financial applications. New Generation Computing pp. 1–15 (2024)
17. Takayanagi, T., Izumi, K., Kato, A., Tsunedomi, N., Abe, Y.: Personalized Stock Recommendation with Investors' Attention and Contextual Information. In: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2023). p. 3339–3343. Association for Computing Machinery (2023). https://doi.org/10.1145/3539618.3591850