InvestAlign: Align LLMs with Investor Decision-Making under Herd Behavior

Paper

Tsinghua University

GitHub

Huisheng Wang, Zhuoshi Pan, Hangjing Zhang, Mingxiao Liu, Yiqing Lin, and H. Vicky Zhao

Background

- LLMs can be leveraged to assist in solving **complex investment problems**.
- However, the investment decisions generated by existing LLMs often **deviate from real-user data**.

How to Align LLMs with Human Decision-Making and Solve Complex Investment Problems with Herd Behavior in the Face of Real-User Data Scarcity?

- In this work, we propose the "**InvestAlign**" framework:

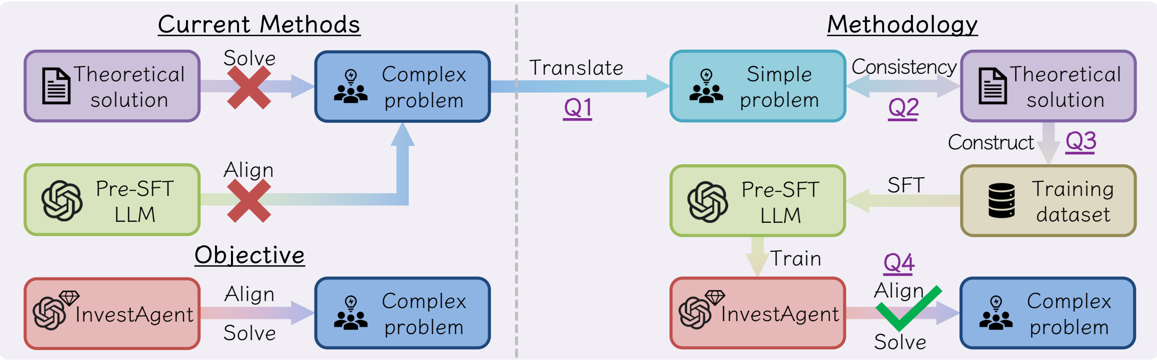

- We need to address four questions, **Q1**, **Q2**, **Q3**, and **Q4**, in the above overview of **InvestAlign** framework.

Q1: Given the complex problem, how do we identify a similar and simpler problem?

Complex problem *P1* (Relative herd): $\sup_{\{P(t)\}_{t\in\mathcal{T}}} \boxed{\mathbb{E}\phi[X(T)]} - \boxed{\theta\delta(P,Q)}$

Translate

Wealth utility   Herd metric

Simple problem *P2* (Absolute herd): $\sup_{\{P(t)\}_{t\in\mathcal{T}}} \boxed{\mathbb{E}\phi[X(T)]} - \boxed{\theta\Delta(P,Q)}$

- **Similar**: *P1* and *P2* are the optimal investment problems with herd behavior and have similar mathematical forms.
- **Simple**: *P2* has an analytical solution of the optimal decision, while *P1* does not have an analytical solution, and calculating its numerical solution is difficult.

Q2: Do the theoretical solution of the simpler problem align with real users' investment decisions, and can they be used to construct a training dataset that mirrors investor decision-making processes?

- We collect real-user data from participants using questionnaires when facing the investment problem *P2*.
- We calculate the difference d and correlation coefficient ρ between the real-user data and the theoretical solution, respectively.
- The t-statistics of d and ρ are -1.075 and -0.843, respectively, both at the 1% significance level.
- We validate that **there exists significant consistency between the theoretical solution and real-user data.**

Q3: How do we generate the training dataset based on the theoretical solution of the simpler problem? How does it perform in aligning with investors' decision-making processes compared with real-user data?

- The SFT training dataset comprises input-output pairs used for fine-tuning LLMs, which are generated based on a custom prompt template.
- When constructing the SFT training dataset, we need to vary the investment attribute, i.e., the risk aversion coefficient α and the influence coefficient θ. We set them through two questions expressed in natural language that are easy for LLMs to understand.

> 🤖 Your **risk aversion coefficient** is {alpha}, which means you consider the following two choices to be indifferent when the probability p is {p}: A. With probability p, you obtain $20, and with probability 1-p, you obtain $0; B. With 100% probability, you obtain $6.
>
> 🤖 Your **influence coefficient** is {theta}, which means in decision-making, your level of dependence on the investment assistant is {k} points. A score of 10 indicates a high level of dependence on the investment assistant, while 0 indicates a low level.

- We theoretically show that fine-tuning LLMs on the training datasets constructed from theoretical solutions leads to **faster parameter convergence** compared to using real-user data and conduct experiments to validate the analysis.

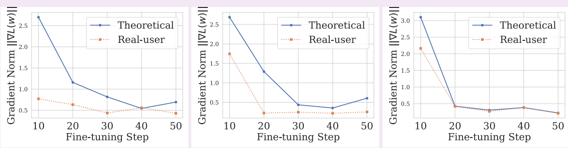

(a) GLM-4   (b) Qwen-2   (c) Llama-3.1

Q4: How do we adapt InvestAgent to solve the complex problem, and what is its performance?

- Performance of **InvestAgent** in *P2*:

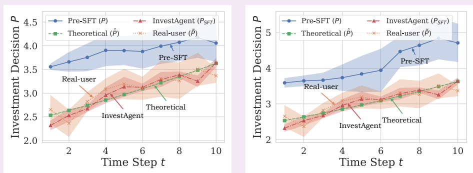

(a) GPT-35   (b) Llama-3.1

- Performance of **InvestAgent** in *P1*:

Table 1: Comparison of the overall MSE between pre-SFT LLMs' and **InvestAgent**s' investment decisions with real-user data in optimal investment problems ***P2*** and ***P1***.

Overall MSE		GPT-35	GLM-4	Qwen-2	Llama-3.1
P2	**Pre-SFT LLM**	4.44	4.20	3.97	4.08
	InvestAgent	1.72	2.26	2.16	1.59
	Reduction from Pre-SFT (%)	-61.26%	-46.19%	-45.59%	-61.03%
P1	**Pre-SFT LLM**	14.03	13.85	17.22	13.07
	InvestAgent	7.46	6.14	7.46	7.25
	Reduction from Pre-SFT (%)	-46.84%	-55.66%	-56.69%	-44.52%

InvestAgent aligns more with humans than pre-SFT LLM in both simple problem *P2* and complex problem *P1*.

Conclusion and Contributions

- LLMs can be leveraged to assist in solving complex investment problems. To fine-tune LLMs for alignment with human decision-making processes, a substantial amount of real-user data is required. However, the cost of collecting the real-user data is high, and there are concerns regarding privacy and security.
- To address these challenges, we propose **InvestAlign**, a novel method that constructs training datasets using the theoretical solution of a similar and simple problem to align LLMs with investor behavior under herd behavior. We demonstrate that fine-tuning LLMs on these training datasets leads to faster parameter convergence compared to using real-user data. The experimental results indicate that **InvestAgent**, fine-tuned with **InvestAlign**, achieves superior alignment performance in the original complex problem.
- Our contributions include: (1) We explore and utilize LLMs in finance and economics, particularly in the domain of optimal investment under herd behavior. (2) We propose the LLM alignment techniques, which construct a large number of high-quality datasets effectively using the theoretical solution of the corresponding mathematical model and then apply SFT to fine-tune LLMs.