Frontiers of Information Technology & Electronic Engineering
www.jzus.zju.edu.cn; engineering.cae.cn; www.springerlink.com
ISSN 2095-9184 (print); ISSN 2095-9230 (online)
E-mail: jzus@zju.edu.cn

Editorial:

Theories and applications of financial large models

Shuoling LIU[†1], Xiaojun ZENG[2], Xiu LI[3], Qiang YANG[†‡4]

[1]*E Fund Management Co., Ltd., Guangzhou 510308, China*

[2]*Department of Computer Science, The University of Manchester, Manchester M139PL, UK*

[3]*Tsinghua Shenzhen International Graduate School, Tsinghua University, Shenzhen 518055, China*

[4]*The Hong Kong University of Science and Technology, Hong Kong 999077, China*

[†]E-mail: liushuoling@efunds.com.cn; qyang@cse.ust.hk

https://doi.org/10.1631/FITEE.2520000

1 Background and motivation

Recent advances in foundation models have ushered in a paradigm shift across the field of artificial intelligence (AI), with profound implications for financial technology (FinTech). Foundation models refer to large-scale neural networks trained on vast and heterogeneous corpora using self-supervised or instruction-driven objectives, which endow them with strong generalization and transfer capabilities across downstream tasks. Representative classes of such models, including large language models (LLMs), multimodal foundation models, and time-series foundation models, exhibit emergent abilities in semantic understanding, reasoning, and multimodal representation learning. These capabilities are fundamentally transforming the operational landscape of financial institutions, including how they process information, evaluate risk, design investment strategies, and interact with clients. Collectively, the rise of foundation models signals a transition toward more adaptive, data-centric, and cognitively informed financial intelligence systems, spanning the entire service lifecycle from risk management and quantitative trading to customer advisory and regulatory compliance.

To further foster this rapidly developing research area and support cross-disciplinary collaboration between AI and finance, *Frontiers of Information Technology and Electronic Engineering* (*FITEE* for short) has organized a Special Feature on "Theories and Applications of Financial Large Models," aiming to solicit original research papers and comprehensive reviews covering both theoretical advancements and practical innovations. The scope of this special feature includes, but is not limited to, the following directions: research on the training, fine-tuning, and evaluation of financial large models (FLMs); generation of intelligent and adaptive portfolio management strategies driven by FLMs; development of quantitative trading algorithms and reinforcement learning frameworks empowered by large models; innovative applications of large models in robo-advisory and personalized financial services; explainability and interpretability of FLMs and their implications for compliance, auditing, and risk management; theoretical and algorithmic studies on the robustness, security, and privacy of FLMs; applications of large models in financial data mining, knowledge discovery, and information extraction; solutions for anomaly detection, fraud analysis, and financial crime prevention incorporating large models; applications of FLMs in environmental, social, and governance (ESG) investing and sustainable finance.

‡ Corresponding author
ORCID: Shuoling LIU, https://orcid.org/0009-0003-1960-3004; Qiang YANG, https://orcid.org/0000-0001-5059-8360

Through these contributions, this special feature seeks to bridge theoretical research with practical deployment, stimulate synergy between academia and industry, and advance the next generation of trustworthy, efficient, and intelligent financial modeling paradigms.

2 Overview of accepted papers

Following a rigorous peer review process, this special feature accepted eight papers contributed by leading academic and industrial institutions around the world. The collection encompasses multiple article types, including one position article, two review articles, three research articles, one correspondence, and one comment. Together, these works provide a comprehensive overview of recent progress in FLMs, ranging from foundational theory and benchmark development to application-oriented financial agent research.

For clarity, the papers are presented under three editorial themes that collectively reflect the field's progression from theoretical innovation to practical deployment. Note that this taxonomy neither implies priority nor mirrors the publication order.

2.1 Foundational frameworks and theoretical advances

The first category centers on the theoretical underpinnings and architectural innovations that establish the conceptual basis for FLMs.

Pushing the boundary of foundational modeling in financial AI, Jian GUO and Heung-Yeung SHUM propose a large investment model (LIM), an upstream foundation model designed to capture cross-market signal patterns from diverse financial datasets. By transferring these globally learned representations to downstream trading strategies, LIM offers a scalable and robust paradigm for universal investment modeling, playing a role analogous to LLMs in natural language understanding. This work exemplifies how foundational modeling principles can be extended from natural language to financial systems.

Further advancing this line of inquiry, Jianzong WANG et al. present a comprehensive survey entitled "*Knowledge distillation for financial large language models: a systematic review of strategies, applications, and evaluation*." Focusing specifically on FinLLMs, they investigate how knowledge distillation can alleviate high deployment costs and long inference latency, offering a practical pathway toward efficient financial intelligence. They structure existing research into three complementary dimensions—strategy, application, and evaluation, and provide a systematic taxonomy for each. In particular, they outline how distillation strategies enhance the practicality of FinLLMs across different stages of the financial workflow, and propose evaluation criteria that account for financial accuracy, reasoning fidelity, and robustness. Overall, this work delivers a balanced and integrative perspective that serves as a roadmap for advancing compact and deployable FinLLMs in real-world financial scenarios.

2.2 Benchmarks, evaluation, and trustworthiness

As FLMs continue to advance, the need for standardized evaluation frameworks and trustworthiness assessments has become increasingly central to ensuring reliability and transparency in real-world deployments.

A growing research trend focuses on reasoning LLMs, which extend general-purpose LLMs by emphasizing multistep inference, logical consistency, and tool-augmented decision making, capabilities that are particularly valuable in financial contexts requiring analytical rigor and interpretability. In this direction, Qiang YANG et al. present one of the first comprehensive examinations of reasoning LLMs in finance. Their study, entitled "*When DeepSeek-R1 meets financial applications: benchmarking, opportunities, and limitations*," benchmarks DeepSeek-R1, a representative reasoning LLM, and its distilled variants on public financial QA datasets. The results provide an early and systematic assessment of how reasoning-oriented models perform in financial applications, revealing both their advantages and limitations. This work offers insights into the emerging role of reasoning LLMs and establishes an early benchmark for assessing their applicability in financial AI scenarios.

Complementing these benchmarking efforts, Shuyan LI et al.'s comment, entitled "*Three trustworthiness challenges in large language model-based financial systems: real-world examples and mitigation strategies*," begins with the observation that integrating LLMs into financial applications introduces unique risks due to the high-stakes nature of financial decision-making. Through a series of finance-specific probing experiments, they identify and analyze three major challenges: jailbreak prompts that exploit weaknesses in model alignment, hallucination leading to factually incorrect or misleading outputs, and bias and fairness concerns that may result in unequal treatment across users or institutions. Based on these findings, they summarize existing mitigation strategies and emphasize that addressing these risks is critical for achieving safe, responsible, and scalable deployment of financial AI systems.

2.3 Practical applications and financial agent systems

The final category highlights application-oriented research that demonstrates how FLMs are being deployed as intelligent agents and analytical tools for decision-making, trading, and forecasting.

In "*FinSphere: a real-time stock analysis agent with instruction-tuned large language models and domain-specific tool integration*," Hongguang LI et al. address two key limitations of current FinLLMs: lack of standardized evaluation metrics for stock analysis and insufficient analytical depth. To tackle these issues, the authors propose AnalyScore, a framework for assessing the quality and interpretability of stock analyses, and construct Stocksis, an expert-curated dataset that strengthens domain-specific modeling capability. Building on these components, FinSphere integrates instruction-tuned financial LLMs with analytical tools to generate professional-grade stock reports. Experiments show that FinSphere consistently surpasses both general-purpose and specialized models in analytical quality and applicability.

Extending the focus from financial analysis to real-world execution, Zhanyu WANG et al. examine whether LLMs can effectively process and execute financial trading instructions, in the paper entitled "*Can large language models effectively process and execute financial trading instructions?*" They design an intelligent trade-order recognition pipeline that converts natural language commands into executable structured instructions. Through evaluation across multiple state-of-the-art LLMs, the work reveals both the potential and limitations of LLM-based trade automation, paving the way for safer human-in-the-loop trading systems.

Moving from trading execution to model-driven financial discovery, Yuchen SHI et al. present a survey entitled "*A survey on large language model-based alpha mining*." The paper explores how LLMs can facilitate hypothesis generation, factor evaluation, and backtesting within an agentic reasoning framework. It identifies key challenges such as limited numerical reasoning, narrow factor diversity, and underexplored factor exploitation mechanisms, while outlining future directions for integrating LLMs as semantic reasoning engines in quantitative research.

Focusing on predictive intelligence and drawing inspiration from narrative economics, Xiaojun ZENG et al. propose a teacher–student multi-agent architecture to detect and forecast evolving events and industry trends, in the paper entitled "*MENTOR: a multi-agent framework for event and narrative trend prediction with optimized reasoning*." Experiments on their self-constructed Chinese key opinion leader articles dataset and English financial news dataset demonstrate consistent improvements in event forecasting, market prediction, and portfolio backtesting. The results highlight how structured reasoning and cooperative agent design can strengthen the link between narrative evolution and market behavior, further expanding the practical frontier of FLMs.

3 Concluding remarks

Collectively, the papers in this special feature provide a panoramic view of the rapidly evolving field of FLMs. They encompass the full research spectrum, spanning foundational theories, benchmark construction, model efficiency, interpretability, and real-world applications. The diversity of topics and contributing institutions underscores a global research ecosystem that unites AI innovation with financial domain expertise.

The guest editors express their sincere gratitude to all authors for their insightful contributions and to the reviewers for their constructive feedback that ensured the quality of this special feature. The guest editors extend their appreciation to the editorial staff of *FITEE* for their continuous support and professionalism. It is our belief that this collection will serve as both a milestone and a catalyst, inspiring future interdisciplinary research on intelligent, interpretable, and trustworthy FLMs.

Shuoling LIU holds a doctorate degree of Computer Science, and is currently the Chief Information Officer of E Fund Management Co., Ltd. He is a member of the FinTech Special Committee of the Asset Management Association of China, a member of the Quantitative Investment Committee of the Investment Technology League (ITL), and a member of IEEE. His research focus lies in the field of financial artificial intelligence (AI), with long-term engagement in advancing technologies and applications in this domain. He has organized several key academic events, serving as the Chairman of the IJCAI 2023/2025 FinLLM Workshop and the ICDM 2024/2025 DMF Workshop. He has won the First Prize once and the Second Prize twice of the FinTech Development Award issued by the People's Bank of China. He also holds over 10 patents related to FinTech innovations.

Xiaojun ZENG is a professor in the Department of Computer Science at the University of Manchester, UK. He received his BS degree in Mathematics and MS degree in Control Theory and Operation Research from Xiamen University, China, and PhD degree in Computation from the University of Manchester, UK. With research experience spanning computational intelligence, machine learning, decision support systems, computational finance, health informatics, and data mining, he has published nearly 200 journal and conference papers; he also led/participated in multiple funded projects (e.g., UK EPSRC, EU Framework Programme, and H2020). He has served the academic community in roles including Associate Editor of *IEEE Trans Fuzzy Syst*, Program Chair of IEEE Symposium on Computational Intelligence in Control and Automation, and Elected Member of the EPSRC College.

Xiu LI is a professor at Tsinghua Shenzhen International Graduate School, Tsinghua University, China. Her research interests include AI, computer vision, and reinforcement learning, with foci in embodied intelligence, multi-modal understanding and generation large models, and AI for life sciences. She has published over 200 academic papers in top-tier international journals and conferences, such as *IEEE Trans Patt Anal Mach Intell*, *IEEE Trans Neur Netw Learn Syst*, CVPR, NeurIPS, and ICML; these papers have been cited over 10 000 times on Google Scholar. She was consecutively listed in the top 2% of global scientists in AI by Stanford University and Elsevier from 2022 to 2025.

Qiang YANG is a Professor Emeritus of the Hong Kong University of Science and Technology, China and the Chief AI Advisor of WeBank. He is a fellow of AAAI, ACM, IEEE, and AAAS, and a member of the Canadian Academy of Engineering and of the Royal Society of Canada. His research interests include transfer learning and federated learning. He is the founding Editor-in-Chief of two journals, *IEEE Trans Big Data* and *ACM Trans Intell Syst Technol*. He has been a conference/program chair for IJCAI, AAAI, and ACM SIGKDD, and has received Distinguished Service Awards from IJCAI and ACM SIGKDD. His latest books are Transfer Learning, Federated Learning, Privacy-Preserving Computing, and Practicing Federated Learning.