Journal of Computing and Electronic Information Management
ISSN: 2413-1660 | Vol. 22, No. 1, 2026

Intelligent Decision Support Driven by AI Agents in Financial Operations

Bokai Wang

School of Electrical Engineering and Computer Science, University of Ottawa, Ottawa, ON K1N 6N5, Canada

Abstract: Artificial intelligence is now used in all parts of a typical company and has changed the form of business. Systematically study the new model of intelligent decision-support systems driven by autonomous, multi-agent AI frameworks in this paper. In the past, because of old, error-prone rule-based computing methods, complex financial risk management and long-term capital allocation were relatively constrained. However, with the emergence of powerful large language models and high-end reinforcement learning algorithms, dynamic and constantly improving AI agents are now being deployed. These intelligent agents cooperate and discuss actively to process large volumes of unstructured market data in sequence, thereby producing highly optimised and perfectly rational financial strategies. The above research indicates that, at present, most models of individual prediction are disconnected from larger-scale, multi-agent automated trading and credit risk systems. In addition, this paper will also investigate how to build an explainable and trustworthy AI system that meets the high standards for international financial regulations. Based on the present study of multi-agent reinforcement learning for algorithmic trading and dynamic risk control, this paper puts forward an all-encompassing system to redesign corporate financial decisions in light of the increasing complexity of the global economy.

Keywords: Artificial intelligence; Financial operations; Multi-agent systems; Decision support; Algorithmic trading; Risk management.

1. Introduction

The scale of the world's financial system is now very large, very high-speed and generates an extremely large amount of data. In the age of high connectivity and severe volatility, the basic level of a company's financial decision-making will determine whether it can be sustained in the long run and how well it will develop. In the past, the many complex financial operations, such as high-frequency algorithmic equity trading, large-scale corporate credit underwriting, and sophisticated capital structure optimisation, were all carried out by traditional rule-based computational models and required continuous human intervention. Although these old systems had a certain degree of computational efficiency, they were also structurally limited. They were inherently static, could not respond promptly to sudden, abnormal macroeconomic shocks, and were very prone to serious cognitive biases and natural fatigue of the human operators.

However, recently, due to an explosive development in artificial intelligence and substantial advances in deep reinforcement learning and complex multi-agent systems, the former limitations have now been overcome. Now that the new financial sector is using AI to make simple forecasts, it is going to build highly autonomous and deep-learning-based AI agents that can conduct extensive and complex financial strategies independently. A single-stage probability score produced by an isolated machine learning model is not the final result of an AI agent; rather, an AI agent can be able to predict changes in its environment proactively, build a multi-stage plan according to this anticipation, carry out specific financial operations in accordance with the plan, and continuously update its knowledge by observing changes in the market.

Robust multi-agent systems (MAS) are increasingly being applied to financial services in a novel way. Several specialised AI agents in the high-end system work together to actively discuss and analyse the financial reasons for different allocation proposals before finally approving a large-scale capital investment plan. The cooperative artificial intelligence system can fully reproduce the complex discussions held by a human executive board, but it operates at a significantly higher speed, mathematical precision and volume of data processing than is feasible for people.

The first aim of this all-encompassing research paper is to systematically study the fundamental alterations these highly intelligent and autonomous AI agents are making to the core structure of financial decision-support systems. Section 2 explores the application of deep reinforcement learning agents in a high-frequency, unstable algorithm trading system. Section 3 presents an application case for multi-agent large language models in complex credit risk assessment and corporate route planning. In Section 4, the reasons for building open and highly reliable AI systems today to meet many demands of financial regulation around the world will be introduced in more detail. Section 5 presents the ending remarks and forecasts the future direction of agent-driven finance. Based on the above advanced technological developments, this paper will investigate how intelligently organised corporate financing will work in the future.

2. Reinforcement learning agents in algorithmic trading

Autonomous AI agents in the area of high-frequency trading are relatively new but may soon bring significant changes to finance [4]. Algorithmic trading fundamentally needs to make very complex and highly constrained capital allocation decisions in a state of extreme market uncertainty [1, 5]. Historically, most quantitative trading strategies have been based on explicit, hard-coded statistical arbitrage rules that are unable to respond promptly to serious market fluctuations and abrupt changes in the distribution of historical prices [2]. Due to the above large-scale operational fragility, more and more financial institutions are now employing high-end reinforcement learning (RL) agents.

Provide a reasonable mathematical basis for the problem of high-frequency trading [4]. In this paradigm, the AI agent is not directly given the list of trades to be carried out. An autonomous agent is continuously interacting with a highly complex simulated financial environment, taking various trading decisions (e.g., buying, holding or short-selling specific assets) at a time, and directly receiving mathematical 'rewards' or 'punishments' based on the profit or loss of its portfolio [5]. After many rounds of quick and intensive training, the deep reinforcement learning agent has developed entirely new high-level trading strategies independently that are far superior to those designed by humans.

In addition, the first of these is that the high-dimensional operational objective space of RL agents can be handled more easily. A good financial agent should not only aim to maximise the absolute raw return but also consider strict constraints such as high transaction costs, changing market liquidity, and the specific, customised risk aversion of the underlying corporate investors [4]. By connecting the basic predictive forecasting problem directly to the overall portfolio construction problem, these intelligent agents have been optimised in a deep machine learning manner to meet the final financial goals of the company. Moreover, a large number of empirical studies have actively demonstrated that when several RL agents are deployed concurrently in highly fragmented and rapidly changing equity markets, they can find deep, hidden statistical arbitrage opportunities long before traditional human analysts are able to process the raw incoming data stream [5, 9].

The application of reinforcement learning agents in financial trading does not guarantee a higher-than-average return unconditionally. Financial markets are non-stationary environments, and at the same time, liquidity, volatility, transaction costs and investor behaviour are all fluctuating over time. Therefore, a good-performing agent in a simulated environment may fail in a real-world change of market conditions. Therefore, the design of trading agents should be subject to strict validation, out-of-sample testing and risk-adjusted performance evaluation. A more practical application of reinforcement learning is an adaptive system for prediction, decision-making and portfolio execution. Based on changes in price or other problems of the market and risk limits, the agent will change its plan. It will be more flexible than a fixed set of trading rules in the system, but it will also need to be monitored by a person. Therefore, reinforcement learning in financial operations should be regarded as a decision-support system for people, rather than a completely independent decision-maker. The first is to handle large volumes of market data, perform many tests on different options, and offer reasonable risk control for these options [1, 4, 5].

3. Multi-agent systems in complex credit assessment

High-frequency algorithmic trading is based on agents for large-scale, high-speed execution, but complex multi-agent systems (MAS) must be employed to address the difficulties in corporate credit risk management and long-term financial planning [10]. Traditional corporate credit scoring used highly opaque and completely rigid logistic regression models to simplify the complex financial situation of large enterprises too much [3, 6]. The older models could not provide in-depth reasons for the serious credit rejection, and thus had to allocate substantial funds unnecessarily [7].

To address this serious structural problem actively, many new financial institutions are using high-end large language model (LLM)-based multi-agent systems to their full capacity [10]. These high-complexity, strongly structured architectures do not use a single, independent AI oracle for the financial decision support system. Instead, the large-scale system deploys a 'chain of agents', which is a highly organised and extremely cooperative network of specialised AIs. For instance, in the evaluation of a large-scale corporate loan application, the top-level 'Analyzer Agent' carefully analyzes a vast amount of complex and unstructured financial data for the company [2, 10]. Afterwards, an extremely critical 'Critique Agent' will systematically list all the initial findings and forcefully identify substantial logical defects, undisclosed leverage risks, or highly unrealistic financial forecasts [10]. Finally, a 'Synthesizer Agent' addresses the previous controversy and proposes an all-round risk-adjusted final credit recommendation.

The first reason for a multi-agent system in credit assessment is that it must generate a final approval or denial. It can be seen that the reasons for a loan decision at different stages are relatively independent. A single agent conducts accounting ratio and cash-flow-stability checks, identifies industry risk and macro-economic exposure, and finally confirms covenant stipulations and repayment capabilities. Division of labour helps to make the decision-making process of a single black-box score more transparent. It can also help the human credit officer identify which part of the assessment has shown a stronger risk signal. With the application of explainable AI in this structure, the system will be able to provide both specific explanations for an individual borrower and general explanations of the primary factors influencing the model. The following reasons are related to the control of credit risk and directly impact the funds and risks of the company. Therefore, the LLM-based or agent-based credit system should be designed to support human review, not to replace it. Organize the large amount of data, identify inconsistencies, and offer supporting analysis for final decisions in finance [2, 6, 7, 10].

The revolutionary debate-based multi-agent framework fundamentally avoids the highly dangerous phenomenon of AI 'hallucination' or severe logical drift [10]. By making the AI agents explicitly justify, fiercely defend, and constantly refine their underlying financial logic through aggressive inter-agent debate, the overall system has significantly increased its absolute reliability, fundamental accuracy, and real-world applicability of the final financial recommendation [6]. These multi-agent systems are also very good at producing rich and consistent natural language explanations for the final financial decisions they have made. All of this transparency has broken the traditional "black box" problem that early machine learning models had, and now human loan officers and corporate executives can use a very understandable, deeply logical, and highly actionable guide for long-term financial optimisation [1, 10].

4. Regulatory compliance and explainable AI frameworks

Although the operationally excellent autonomous AI agents are of significant benefit, their full-scale and unrestricted use in the financial sector around the world carries substantial, very complex regulatory and systemic risks [8-9]. The world's financial system has been subject to the construction of a very strict and highly inflexible regulatory system that strictly aims to avoid large-scale systemic risks, severe market misconduct and deep-seated discriminatory lending practices [3]. As these very intelligent AI systems have been developing at an accelerated pace and are moving from simply offering passive advice to operating as fully autonomous financial agents deeply integrated with large-scale payment infrastructures and real-world capital, the basic concept of 'trust' is changing drastically, moving from a statistical measure of correctness to one of absolute end-to-end product security [7, 9].

Financial regulators now demand complete transparency without reservation [3]. If a highly complex multi-agent system suddenly refuses to approve a large-scale corporate credit line or rapidly sells off a substantial block of shares, the financial institution must be able to fully explain the exact, highly specific mathematical and logical reasons for such a decision [7]. This strong, non-negotiable regulatory demand has led to an all-out, fast-paced expansion of Explainable AI (XAI) in contemporary financial risk management [8]. Explainable AI frameworks are explicitly designed to completely demystify the extremely complex inner workings of advanced neural networks, ensure that the final decisions of the AI agent are in full accordance with strict ethical standards, are absolutely free from hidden historical biases, and comply with rigid international banking laws [7, 9].

Researchers and financial engineers have also been building entirely new, very strong risk-management systems that can be used to audit and strictly supervise the random-walk behaviour of these autonomous agents [8]. Given that agentic behaviour is inherently dynamic and highly unpredictable, the traditional static software testing model is no longer feasible [9]. Financial institutions should now use large-scale, persistent adversarial tests to have specialized 'Red Team' AI agents proactively try to deceive, mislead or force the main financial agents into executing serious and unlawful financial activities [7]. Only by subjecting these multi-agent systems to the most severe and harsh simulated market conditions can human executives be assured that the AI agents will perform their massive financial duties reliably, safely, and highly ethically in the real world [8-9]. In the end, absolute trustworthiness and strong explainability are not only good philosophical ideas but will also serve as the foundation for the entire future of automated agentic finance and need to be built into the system.

5. Conclusion

A large-scale, entirely new model of global corporate operations will be implemented by integrating autonomous and highly intelligent AI agents deeply into the core of financial decision-support systems. As this all-encompassing analysis has shown, the traditional era of using only rigid, highly static predictive models and deeply flawed human intuition is rapidly coming to an end. Aggressively deploy high-end deep reinforcement learning algorithms and sophisticated, debate-driven multi-agent language models to implement large-scale, highly dynamic trading and credit strategies at the speed and absolute mathematical accuracy far surpassing what humans can achieve naturally.

Reinforcement learning agents have been able to discover some unknown, very high-profit statistical arbitrage opportunities in the extremely volatile environment of algorithmic trading while managing transaction costs and dynamic risks effectively. At the same time, in light of the above, together with this, cooperative multi-agent systems have also been actively developed and studied to address the very difficult problem of corporate credit allocation; such systems can process unstructured financial data to generate excellent, perfectly rational and extremely transparent capital recommendation plans. With good results achieved so far, many high-tech solutions have improved both the overall efficiency and foundational strength of contemporary finance.

However, the final results of this large-scale technological revolution will depend on how well the financial system addresses the serious and inherent systemic risks caused by absolute autonomy at the level of detail. Given that these highly intelligent agents will have full and unsupervised control over vast amounts of global capital, the complete application of strict, fully explainable and deeply trustworthy AI governance systems is now unavoidable. Proactively and in full cooperation with foreign regulatory authorities, establish strict standards for agentic risk underwriting and carry out continuous adversarial auditing at the highest level. By successfully connecting high-performance computing with absolute structural reliability, the entire global financial sector will be enabled to realise its full potential of AI agents while ensuring the long-term safety, unwavering fairness and foundational stability of the world economy.

References

[1] Sutton, R. S., & Barto, A. G. (2018). Reinforcement learning: An introduction (2nd ed.). MIT Press.

[2] Healy, P. M., & Palepu, K. G. (2001). Information asymmetry, corporate disclosure, and the capital markets: A review of the empirical disclosure literature. Journal of Accounting and Economics, 31(1-3), 405-440. https://doi.org/10.1016/S0165-4101(01)00018-0

[3] Leuz, C., & Wysocki, P. D. (2016). The economics of disclosure and financial reporting regulation: Evidence and suggestions for future research. Journal of Accounting Research, 54(2), 525-622.

https://doi.org/10.1111/1475-679X.12115

[4] Fischer, T. G. (2018). Reinforcement learning in financial markets: A survey. FAU Discussion Papers in Economics, 12, 1-35.

[5] Liu, X. Y., Yang, H., Chen, Q., Zhang, R., Yang, L., Xiao, B., & Wang, C. D. (2021). FinRL: Deep reinforcement learning framework to automate trading in quantitative finance. In Proceedings of the Second ACM International Conference on AI in Finance (pp. 1-9). ACM.

https://doi.org/10.1145/3490354.3490373

[6] Misheva, B. H., Osterrieder, J., Hirsa, A., Kulkarni, O., & Lin, S. F. (2021). Explainable AI in credit risk management. arXiv, arXiv:2103.00949. https://doi.org/10.48550/arXiv.2103.00949

[7] Fritz-Morgenthal, S., Hein, B., & Papenbrock, J. (2022). Financial risk management and explainable, trustworthy, responsible AI. Frontiers in Artificial Intelligence, 5, Article 779799. https://doi.org/10.3389/frai.2022.779799

[8] Giudici, P., Centurelli, M., & Turchetta, S. (2024). Artificial intelligence risk measurement. Expert Systems with Applications, 235, Article 121220.

https://doi.org/10.1016/j.eswa.2023.121220

[9] Vyas, A. (2025). Revolutionizing risk: The role of artificial intelligence in financial risk management, forecasting, and global implementation. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.5012345

[10] Yarmohammadtoosky, S. (2025). Enhancing credit path planning with LLM-based multi-agent systems [Doctoral dissertation, Kennesaw State University].