Evolution of machine learning in financial risk management: A survey

Kuan-I Lu

Actuarial Science, Department of Applied Probability and Statistics, 93117 University of California, Santa Barbara, United States

Abstract. Financial risk management plays a crucial role in daily financial decision-making, aiming to mitigate risk and maximize profit. Given its reliance on data, financial risk management can greatly benefit from the application of machine learning tools. Over the years, we've observed a clear trend in the evolution of these applications, marked by increasing model complexity and a broader range of manageable tasks. This paper contributes to the field in three key dimensions: First, we provide a clear taxonomy of risks and an introduction to relevant machine learning methods to establish a foundation and identify the targeted issues. Next, we explore real-world data applications, discussing the pros and cons of three methods, from the earliest to the most recent. Finally, based on the observed results, we highlight current challenges and limitations and propose potential directions for improvement.

1 Introduction

Machine Learning (ML) has long played a crucial role in the era of artificial intelligence. From prediction tasks to signal and image processing, and more recently to breakthroughs in natural language processing, the impact and potential of ML have become increasingly evident. Its applications in Financial Risk Management (FRM) are no exception, with its effectiveness well-validated [1]. Whether forecasting future Value at Risk (VaR) trends based on historical data, predicting customer default rates, or managing financial portfolios, ML proves especially valuable in the financial sector, where data is abundant and the need for improved performance drives the adoption of data-driven approaches.

Many previous surveys have offered comprehensive and unified overviews of ML applications in FRM tasks [2, 3]. However, this paper approaches the issue from a different angle and aims to provide an overview of the evolution of applications in a chronological order. We illustrate the progression from simply presuming a time series model and then refine it with Bayesian inference, to using more complicated structures like neural networks to mitigate the inherent complexity of financial data, to elevating the task form metric prediction to decision making with the use of reinforcement learning (RL). Given the high-stakes nature of FRM, even small mistakes or false assumptions can result in significant

Corresponding author: 1k_lu@ucsb.edu

financial losses. Consequently, stakeholders often prefer explainable models, which allow for clearer interpretation of the model's logic, as discussed in [4]. However, despite this preference for simpler, transparent methods, more complex black-box models, which have demonstrated significantly greater effectiveness, seem inevitable for accurate predictions and the full utilization of ML in FRM. This paper explores this transformation and supports it with evidence from related works.

In this paper, we begin by breaking down the broader term "Financial Risk Management" into several classifications. This taxonomy helps clarify the purpose of each ML model and assess the versatility of each method based on the types of risk it can address. By using this taxonomy, we can better define the scope of the problem and more easily understand the characteristics of the targeted task. Additionally, we can share knowledge directly with those facing the same types of risks. Next, we provide a brief background on the ML methods discussed in this overview, ranging from the simplest Bayesian Inference to the more complex Neural Networks, and finally to RL, which advances the task from value prediction to decision-making. The sequence of these methods illustrates not only increasing model complexity but also improved performance on both experimental and real-world data. Finally, we highlight the inherent flaws, examine the current challenges these models face, and propose potential directions for future improvements. Our aim is to help financial risk managers integrate ML into real-world financial settings, making their jobs easier, enhancing performance, and ultimately generating more revenue through more effective risk assessment.

The goal of this paper is not only to provide a high-level overview to help readers, even those without deep expertise in FRM, to easily grasp the power of ML models, but also to demonstrate and explain their effectiveness. The paper delves into each method, presenting numerical data from related research. It serves as both a handbook for those unfamiliar with ML and a guide for deepening their knowledge once they have gained a general understanding of these tools and their applications.

The organization of this overview is as follows: Section 2 presents a taxonomy of financial risks by discussing four distinct types and providing a brief introduction to each corresponding ML method. Section 3 explores the application and performance of Bayesian inference, with its effectiveness verified through prior experiments. Section 4 takes on a similar task, focusing on neural networks, while Section 5 does the same for RL. Finally, Section 6 highlights the challenges current methods face and suggests potential future improvements to enhance these methods and their applications.

2 Classifications of financial risks and ML methods

2.1 Financial Risks

Before delving into the applications and solutions, it is essential to first identify the different types of financial risks and examine the information they reveal. A basic taxonomy of these risks is outlined in [2, 5], where financial risk is categorized into four main types:

- *Market risk* refers to external factors affecting the market, such as stock prices, interest rates, and exchange rates. Common measures of market risk include VaR, which estimates the maximum potential loss of a portfolio over a specified period; Conditional Value at Risk (CVaR), a more comprehensive version of VaR that accounts for tail risk; and Systemic Risk (SRisk), which estimates a firm's capital shortfall during a financial crisis.
- *Credit risk* refers to the risk that either the borrower or the lender fails to meet their obligations, potentially resulting in financial loss for either party. Common measures of credit risk include Default Risk, which assesses a borrower's credibility and the

likelihood of loan repayment failure, and Counterparty Risk, which gauges the probability that the counterparty in a financial transaction will fail to fulfill their obligations, particularly in trading and investment activities.

- *Operational risk*, unlike market risk, refers to the internal risks that arise from a company's operations. Different firms may have distinct Key Risk Indicators (KRIs), such as power outage frequency, reported fraud incidents, and manager turnover rates, among others. These metrics can be categorized into business risks and event risks, with business risks focusing on routine operational challenges, and event risks addressing significant internal events that could lead to substantial losses.
- *Insurance and demographic risks* are most observed in the insurance industry, where companies focus on events that may require large claim payouts, known as Catastrophe Risk, or on changes in the population's overall mortality rate, referred to as Mortality Risk. The mitigation of these risks is typically reflected in the pricing of insurance products, but since that is not the primary focus of this paper, it will not be discussed further.

In this paper, we focus on market risk and credit risk, two of the most widely discussed and applicable topics in modern FRM research. Market risk data is typically time-dependent and non-linear, with the goal of identifying trends, as well as the correlations and causations among key indicators of the overall financial environment. For credit risk, the primary task is to assess and estimate a borrower's or counterparty's creditworthiness based on their portfolio and other relevant factors.

Estimating these risks is not straightforward. Relying solely on mathematical methods may overlook subtle underlying mechanisms, while simulations might overgeneralize data and miss extreme cases. For example, using time series models to forecast future risk trends can be inefficient, as observed data is often too sparse to predict major crises or extreme events. This highlights the need for more sophisticated methods for risk prediction and estimation.

2.2 ML Methods

The basic classifications of ML methods are described in [6]. Rather than offering a comprehensive overview of all ML methods, relevant or not, this section focuses specifically on the methods discussed in this paper, presented in chronological order. Each subsequent method addresses the flaws and weaknesses of the previous one. The methods covered are *Bayesian Inference, Neural Networks,* and *Reinforcement Learning*, and a brief introduction and explanation of each are provided below:

- *Bayesian Inference* [7] requires an initial assumption of parameters for a given model—in this case, a time series model explaining the trend and fluctuation of market risk measures. The model is then updated with observed data. This approach primarily focuses on explaining time series data, such as market risk metrics like VaR and CVaR.
- *Neural Networks* [8] mimic the structure of the human brain, passing compressed information through neurons in successive layers to produce an output. A loss function is used to adjust the model's weights via backpropagation, refining it to best explain the data. Neural networks' flexibility and high accuracy make them suitable for handling various types of financial risk data, leading to their widespread use in modern financial risk management.
- *Reinforcement Learning* [9] models an agent interacting with its environment. Through reward feedback at each state, the agent learns the optimal strategy and actions to maximize overall reward. This method bypasses the need for predicting metrics, focusing instead on determining the next best action—an approach aligned with the goals of investors and firms.

The next three sections explore the application of each method using real-world data, accompanied by critiques and suggestions based on their performance in each case. Additionally, a simplified workflow diagram for each method is included to enhance understanding of their significance and practical value.

3 Bayesian inference

3.1 Preliminary

In contrast to the widely used frequentist approach, which directly uses observed data as an estimator of the true value, Bayesian inference is more effective when dealing with limited or sparse data, as is often the case with financial datasets. Bayesian inference constructs a prior distribution—representing prior knowledge or "estimation from experience"—and assigns a degree of certainty, or weight, to this estimate. In this framework, the data's role is to update the prior estimate, adjusting it based on the observed information. This method is particularly useful in financial statistics, where accessible data is often sparse, and conclusions drawn solely from observations can have high variance, leading to inaccurate results and predictions. The main advantage of Bayesian inference is that it introduces some bias to reduce overall error, balancing the bias-variance tradeoff.

3.2 Applications and Experiments

Being one of the earliest approaches to assess risk by constructing a time series model to explain market risk, Bayesian inference is significant for its implementation of human knowledge and generalization. As we know, financial risk data, particularly VaR data, is time dependent, multivariate, and nonlinear. The best way to explain the data with a model would be using a time series model like DCC-GARCH, DCC-GJR-GARCH, IGARCH, EGARCH, Markov switching GARCH, just to name a few [10, 11], while the forecasting and estimation employs Monte Carlo Markov Chain (MCMC) [11]. These are models that handle multivariate dependencies, allowing correlations between variables to change over time, built on top of the classic GARCH model that explains financial data with varying variances. The part Bayesian inferences play here is to, by using prior domain knowledge and experience, identify the vectors of AR(s) coefficients, Φ, and the vector of volatility coefficients, α. Then we update the coefficients, which dictates the model defining the trend and prediction of the financial data.

An experiment demonstrating the usefulness of Bayesian inference in financial risk analysis examined high-frequency data from ten component stocks of the DJIA index in 2010, using eight variations of the GARCH model [10]. By focusing on stock prices, the study primarily addressed market risk. In another example, researchers analyzed four different stock markets, fitting multiple GARCH models with a Bayesian approach, and achieved similarly satisfactory results [11]. The general workflow of Bayesian inference in risk management is illustrated in Fig. 1, which is straightforward. Researchers begin with a prior distribution, whether informative or not, and then update it with observed data. The more complex task lies in selecting and tuning the best model to describe financial time series data and pre-selecting the parameters.

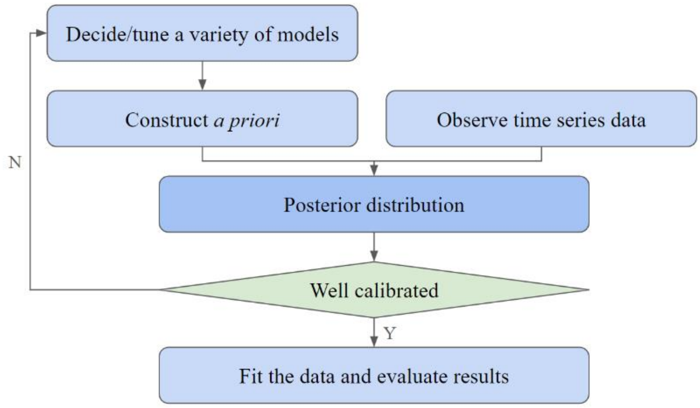

Fig. 1. Workflow of constructing model explaining financial risk data with Bayesian inference**.**
(Photo/Picture credit: Original)

3.3 Performance

The most valuable feature of Bayesian inference is also its greatest flaw, by allowing to initiate the selected model with an a priori, we can make sure our model is not fully dependent on the objective data, which can be misleading due to its sparsity and could lead to more inaccurate results. However, Bayesian inference is only useful when the proposed prior distribution is a good estimate of the real distribution/value. If one starts off with an extremely biased prior, relying solely on the objective data may be a better approach.

Furthermore, the Bayesian inference approach remains grounded in the mechanistic model-fitting method. By assuming the VaR data follows a GARCH model, we limit the flexibility and potential of exploring other models or explanations for the data, which can lead to overgeneralization. Even when we establish a strong a priori, the model choice itself may still contribute to inaccurate results. While Bayesian inference, paired with the wide variety of GARCH models available for tuning, has generally performed well, there remains significant room for improvement. Its limitations stem from its mechanistic nature and its inability to accurately predict major events, as highlighted by the disparity in performance before and after the financial crisis, with post-crisis forecasts being less accurate and overly conservative [11].

4 Neural Network

4.1 Preliminary

As a more complex model in the evolution of ML, the neural network is structured similarly to the human brain, with nodes functioning like neurons. The most basic form of a neural network is the Artificial Neural Network (ANN), where the input variables are first normalized and encoded into an input layer. Each node in this layer is then processed with weights and activation functions, feeding into the nodes of the next "fully connected" layer, known as a dense layer. This process is repeated across subsequent dense layers, depending on the model's complexity. Ultimately, the dense layer is connected to the output layer, with the number of nodes determined by the task at hand. For example, in binary classification, the output layer would contain a single node indicating the probability of one of the two classes. During model fitting, the neural network calculates a loss function based on its predictions and uses backpropagation or other learning methods to adjust the weights and optimize the parameters [13]. A visual representation of a typical ANN structure can be found in Fig. 2.

Variations of the vanilla ANN include: Deep Neural Networks (DNN), which consist of multiple dense layers [14]; Convolutional Neural Networks (CNN), where only a portion of nodes are connected to the next layer, making them particularly effective for image processing; Recurrent Neural Networks (RNN), which use recurrent structures to better handle sequential data; and Long Short-Term Memory (LSTM) networks, a type of RNN designed to retain important information over long time periods [15].

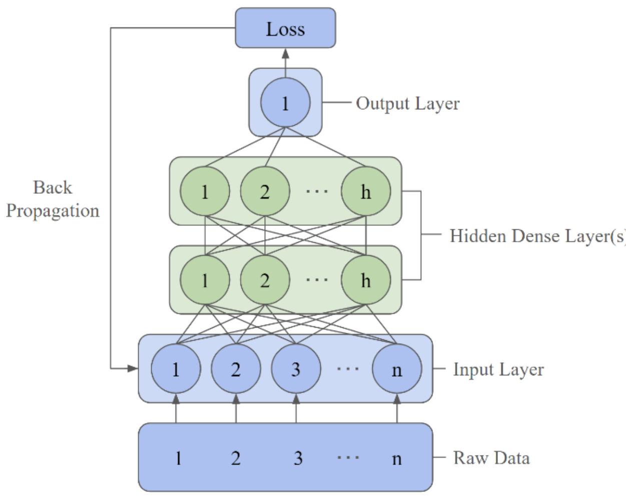

Fig. 2. General workflow of an ANN model training structure. (Photo/Picture credit: Original)

4.2 Applications and Experiments

Neural networks are considered empirically driven models rather than mechanistic ones, giving them strength in versatility and broad applicability across various risk management tasks. A supervised neural network can handle functions such as predicting default rates, assessing peer-to-peer lending risk, credit scoring, and market forecasting by processing numeric data. Additionally, they can perform fraud detection by processing natural language, aligning well with the broader scope of financial risk management. Other applications include portfolio management using unsupervised neural networks and stock market analysis with LSTM [16].

In managing credit risk, earlier research demonstrated the effectiveness of neural networks in credit risk evaluation and prediction by comparing several ANN models with varying levels of complexity. This experiment, using a German credit dataset, achieved a validation accuracy of up to 73% [17]. More recent studies explored the application of CNNs in predicting the financial risk of Chinese companies, with the best models reaching an accuracy of approximately 81%. This research also applied PCA to reduce parameters and training time, significantly cutting the typically high computational cost of neural networks [15]. Other researchers employed more complex DNN architectures, including ANNs with multiple hidden layers and larger datasets, achieving an overall accuracy of 93%, well above the benchmark of 73% [16]. The latest advancements involve training neural networks on large datasets generated by the Internet of Things (IoT), using a PSO-based BP neural network, which has achieved an impressive precision of 99% [18].

4.3 Performance

Although the Bayesian Inference approach did not provide quantitative results in terms of error or precision, neural network approaches, regardless of the version, demonstrated improvements in both versatility and performance when handling financial data and predicting financial risk. As the most widely used ML method for FRM tasks, neural networks are the most well-researched and mature, consistently delivering reliable results. These models not only predict time-independent metrics, such as default rates, but also handle time-dependent data, like market risk values, where time series models were traditionally applied. The evolution of various neural network models has further enhanced the robustness of this approach, making it capable of processing and explaining a wide range of data types.

While neural networks have proven to be effective and reliable for prediction tasks, there remains a significant flaw. The gap between accurately predicting a key index and making the appropriate decisions based on that prediction still relies heavily on human judgment. None of the ML methods previously mentioned have fully addressed this issue. In other words, understanding risk-related values and knowing how to act on them are two entirely different objectives, with the latter being more critical. Addressing this challenge requires a fresh perspective, which we will explore in the next section.

5 Reinforcement Learning

5.1 Preliminary

RL is a branch of ML that differs from supervised learning, where an outcome variable serves as a guide during the training process. In contrast, RL employs an agent-based approach, focusing on situations that require making long-term, optimal decisions within a given environment. The goal is to maximize the cumulative rewards earned over time. Through feedback in the form of rewards or punishments, the agent refines its policy—a mapping from states to actions—via trial and error, to achieve long-term objectives [19]. In complex systems like the financial market, RL is particularly powerful due to its ability to adapt to unpredictable conditions and learn strategies that optimize long-term outcomes, even amidst short-term uncertainty.

To better illustrate the workflow of RL, Fig. 3 shows how the environment provides feedback to the agent in the form of a reward, allowing it to update its policy. The agent's policy then determines which action to take in each state, aiming to maximize the accumulated reward.

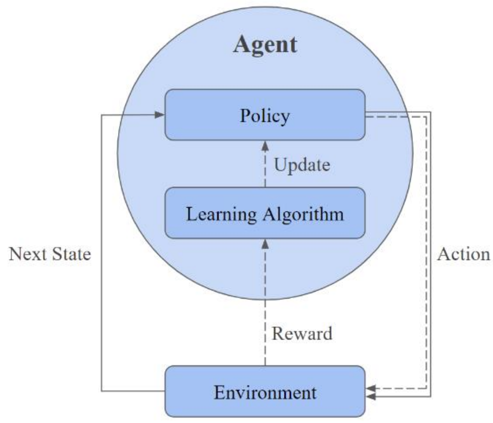

Fig. 3. General workflow of RL. (Photo/Picture credit: Original)

5.2 Applications and Experiments

Instead of assessing risk and predicting numerical values, RL directly focuses on learning the best strategies and actions to achieve maximum long-term rewards. This approach differs from traditional supervised or unsupervised ML tasks, which typically aim to predict or infer specific values or indices. In RL, risk control measures can range from managing market dynamics and credit risk to optimizing the portfolio itself. In other words, strategies developed through RL can uncover patterns and tactics that may not be obvious or intuitive. By bypassing the price or index prediction typically required in Bayesian inference or neural network models—where an additional hand-coded layer is needed to convert predictions into actionable strategies—RL simplifies the process by targeting optimal actions directly from the outset [20].

One of the most significant applications of RL in financial risk tasks is portfolio allocation strategy. In one study, a researcher developed a RL framework to manage cryptocurrency portfolios and handle multi-channel market inputs. The output portfolio weights directly influenced market actions, and the framework was integrated into several deep neural networks, including DNN, RNN, and LSTM. The model is also linearly scalable with portfolio size [20]. This approach represents an advancement over previous applications of ML in cryptocurrency [21]. Another study used rank-dependent expected utility (RDEU) as a risk measure. Researchers applied a RL model based on the results of a neural network with three hidden layers and 50 neurons. This model actively mitigated risk by adjusting the distribution of terminal wealth, shifting lower terminal values to mid-level compared to the benchmark strategy [21].

5.3 Performance

From the previous examples, it's clear that RL does not replace the neural network approach. Instead, it serves as a next-step methodology for training machine-learned decision-making processes, demonstrating its robustness in refining strategies based on actions generated by neural network outputs. However, its limitations arise in its inability to handle a broader range of tasks. In the examples discussed, the primary applications are portfolio management and algorithmic trading, where RL naturally excels. That said, there is potential for RL to be applied to tasks beyond trading, with the possibility of yielding similarly strong results.

6 Current challenges and future directions

The exploration demonstrated the proficiency of these models but also revealed flaws and areas for improvement. Bayesian Inference, our earliest approach, serves as the baseline method, with subsequent methods compared to its performance. While the Bayesian approach's predictions are reasonable, they remain limited by the method's narrow application, focusing solely on time series forecasting, which is insufficient for capturing the complexity and variety of financial risk measures. Moreover, Bayesian inference requires a pre-assumed model and parameters, leaving too much room for human interference. Unless one can miraculously identify the correct underlying model that explains the data, the bias introduced—while it may reduce overall loss—still leaves room for improvement that other methods can more effectively address.

Improvements in accuracy are clear when using neural networks. As one of the most popular and well-developed areas of modern research, neural network variations have performed exceptionally well in predicting correct values when provided with sufficient data and computational power. However, one challenge is the need for an adequate amount of training data. The biggest issue with neural networks is the accessibility of relevant data. For metrics that naturally lack sufficient data, neural networks can lead to overfitting, where a more mechanistic approach may outperform them.

The shift from neural networks to RL offers a new perspective on solving FRM tasks. Rather than predicting a metric and developing strategies around that prediction, RL directly identifies the optimal move by learning the underlying strategy, with the goal of maximizing rewards—typically profit. As RL is still in the early stages of application to FRM tasks, several key areas for improvement are evident. First, the simplicity of the reward function: decision-makers often consider factors beyond profit, and a model that focuses solely on profit may not fully align with human intentions. Second, current use cases for RL are largely confined to investment portfolio management. While these applications have shown promising results, there is a strong need to explore other areas, such as loan distribution based on credit risk [23] or adjustments to economic policies in response to systemic risks [24]. Strategizing requires deep contemplation and substantial experience, but if future advancements in RL can replicate these complex processes, it could significantly reduce decision-making costs, allowing for greater investment of capital. Though RL's application in FRM is still in its early stages, the potential for machine-learned strategies to assist in the strategic process is promising and could be groundbreaking in the future.

Concerns about these methods may arise, as neural networks and RL are often considered black-box approaches, making them difficult to explain. Additionally, issues such as data accessibility, cybersecurity in the implementation of these methods, and the availability of skilled ML professionals could present challenges when deploying these techniques in the market. Moreover, this paper primarily addresses market and credit risk, leaving out two other categories: insurance risk and operational risk. Future improvements could focus on accurately predicting these types of risk using alternative ML methods, despite challenges like dataset sparsity. Training a high-performing model on such inefficient data is a complex task, and I look forward to seeing future proposals tackle these risks.

7 Conclusion

Artificial Intelligence, particularly ML, has become an unstoppable force in today's world. As individuals and organizations strive to incorporate these technologies into their daily tasks, a chronological overview of how these methods have evolved can help identify areas for future improvement. In this paper, we successfully analyzed the progression and enhancement of techniques from Bayesian Inference to Neural Networks, and finally to RL, focusing on their capacity to manage FRM over time. We demonstrated that while Bayesian Inference has its uses, it is the most limited in terms of practical applications. Neural Networks, despite being primarily focused on predicting related risk values, have proven to be highly versatile, with widespread industrial adoption and demonstrated effectiveness. RL, though still in its early stages of development, shows great potential as a decision-making tool that can build upon the results generated by neural networks. As suggested in [1], although current methods have already revolutionized how FRM tasks are approached, there remains significant room for improvements. These improvements range from aspects as specific as the exploration of application on other tasks using RL, to broader issues such as data accessibility, security, talent, and the mitigation of other types of financial risks. Overall, the trends and insights presented in this paper should positively influence those in FRM seeking to drive innovation and enhance performance.

References

1. N. Milojević, S. Redzepagic, Prospects of artificial intelligence and machine learning application in banking risk management. J. Cent. Bank. Theory Pract. 10.3, 41-57 (2021). https://doi.org/10.2478/jcbtp-2021-0023
2. A. Mashrur, W. Luo, N. A. Zaidi, A. Robles-Kelly, Machine learning for financial risk management: A survey. IEEE Access. 8, 203203–203223 (2020). https://doi.org/10.1109/ACCESS.2020.3036322
3. M. Leo, S. Sharma, K. Maddulety, Machine learning in banking risk management: A literature review. Risks, 7, 29 (2019). https://doi.org/10.3390/risks7010029
4. N. Bussmann, P. Giudici, D. Marinelli, J. Papenbrock, Explainable machine learning in credit risk management. Comput. Econ. 57, 203-216 (2021). https://doi.org/10.1007/s10614-020-10042-0
5. J. Mulvey, D. Rosenbaum, B. Shetty, Strategic financial risk management and operations research. Eur. J. Oper. Res. 97, 1–16 (1997). https://doi.org/10.1016/S0377-2217(96)00222-6
6. S. Aziz, M. Dowling, AI and machine learning for risk management. Disrupting Finance, 33–50 (2019). https://doi.org/10.2139/ssrn.3201337
7. G. D'Agostini, Bayesian inference in processing experimental data: principles and basic applications. Rep. Prog. Phys. 66, 1383-1420 (2003). https://doi.org/10.1088/0034-4885/66/9/201
8. D. E. Rumelhart, B. Widrow, M. A. Lehr, The basic ideas in neural networks. Communications of the ACM. 37.3. 175247-175256 (1994). https://doi.org/10.1145/175247.175256
9. L. P. Kaelbling, M. L. Littman, A. W. Moore, Reinforcement learning: A survey. J. Artif. Intell. Res. 4, 237-285 (1996). https://doi.org/10.1613/jair.301
10. E. M. H. Lin, E. W. Sun, M.-T. Yu, Behavioral data-driven analysis with Bayesian method for risk management of financial services. Int. J. Prod. Econ. 228, 107737 (2020). https://doi.org/10.1016/j.ijpe.2020.107737
11. C. W. S. Chen, R. H. Gerlach, E. M. H. Lin, W. C. W. Lee, Bayesian forecasting for financial risk management, pre and post the global financial crisis. J. Forecast. 34.8, 661-687 (2012). https://doi.org/10.1002/for.1237
12. Q. Feng, H. Chen, R. Jiang, Analysis of early warning of corporate financial risk via deep learning artificial neural network. Microprocess. Microsyst. 82, 104387 (2021). https://doi.org/10.1016/j.micpro.2021.104387

13. X. Li, J. Wang, C. Yang, Risk prediction in financial management of listed companies based on optimized BP neural network under digital economy. Neural Comput. Appl. 35, 2045-2058 (2023). https://doi.org/10.1007/s00521-022-07377-0
14. A. Kim, Y. Yang, S. Lessmann, T. Ma, M. C. Sung, J. E. V. Johnson, Can deep learning predict risky retail investors? A case study in financial risk behavior forecasting. Eur. J. Oper. Res. 283.1, 217-234 (2020). https://doi.org/10.1016/j.ejor.2019.11.007
15. B. Wang, A financial risk identification model based on artificial intelligence. Wirel. Netw. 30, 4157-4165 (2024). https://doi.org/10.1007/s11276-021-02856-z
16. J. Duan, Financial system modeling using deep neural networks (DNNs) for effective risk assessment and prediction. J. Franklin Inst. 356.8, 4716-4731 (2019). https://doi.org/10.1016/j.jfranklin.2019.01.046
17. A. Khashman, Neural networks for credit risk evaluation: Investigation of different neural models and learning schemes. Expert Syst. Appl. 37.9, 6233-6239 (2010). https://doi.org/10.1016/j.eswa.2010.02.101
18. H. Zhou, G. Sun, S. Fu, J. Liu, X. Zhou, J. Zhou, A big data mining approach of PSO-based BP neural network for financial risk management with IoT. IEEE Access. 7, 154035-154043 (2019).zhttps://doi.org/10.1109/ACCESS.2019.2948949
19. A. Charpentier, R. Elie, C. Remlinger, Reinforcement learning in economics and finance. Comput. Econ. 62, 425-462 (2023). https://doi.org/10.1007/s10614-021-10119-4
20. Z. Jiang, D. Xu, J. Liang, A deep reinforcement learning framework for the financial portfolio management problem. arXiv, 1706.10059 (2017). https://doi.org/10.48550/arXiv.1706.10059
21. Z. Shahbazi, Y. C. Byun, Machine learning-based analysis of cryptocurrency market financial risk management. IEEE Access, 10, 37848-37856 (2022). https://doi.org/10.1109/ACCESS.2022.3162858
22. S. Jaimungal, S. M. Pesenti, Y. S. Wang, H. Tatsat, Robust risk-aware reinforcement learning. SIAM J. Financial Mathematics, 12.4, 1234-1256 (2021). https://doi.org/10.48550/arXiv.2108.10403
23. D. Mhlanga, Financial inclusion in emerging economies: The application of machine learning and artificial intelligence in credit risk assessment. Int. J. Financial Stud. 9.3, 39 (2021). https://doi.org/10.3390/ijfs9030039
24. M. Stolbov, M. Shchepeleva, Systemic risk, economic policy uncertainty and firm bankruptcies: Evidence from multivariate causal inference. SSRN Electron. J. 36, 1-36 (2019). https://doi.org/10.2139/ssrn.3502420