Journal of Computer Technology and Software
ISSN:2998-2383
Vol. 3, No. 8, 2024

Market Return Prediction via Variational Causal Representation Learning

Yuan Sheng
Northeastern University, Seattle, USA
Sophia.sheng.ms@gmail.com

Abstract: This paper proposes a multi-factor market return prediction model based on variational causal representation learning. The goal is to improve prediction accuracy, robustness, and interpretability. The method adopts a variational inference framework to learn latent causal structures from high-dimensional factor data. It incorporates a causal regularization term and a counterfactual consistency loss to enhance the model's resistance to spurious correlations and data perturbations. The model consists of three components: a variational encoder, a generator, and a predictor. It is trained end-to-end to jointly learn causal representations and perform return prediction. In multiple experiments, the proposed method outperforms representative existing models in terms of mean squared error, causal consistency, and robustness. It also shows strong adaptability in transfer learning tasks across different economic cycles. Ablation studies confirm the contribution of each module to overall performance. These results further demonstrate the value of causal modeling in improving the stability of financial prediction models.

Keywords: Causal inference, variational inference, representation learning, market return prediction

1. Introduction

The financial market, as a vital part of the economic system, operates under complex and highly dynamic patterns. It is influenced by a variety of macroeconomic factors, industry changes, and micro-level market behaviors. For a long time, market return prediction has been a central issue in finance, attracting significant attention from both academia and industry. Traditional multi-factor models have, to some extent, revealed the major drivers of return fluctuations [1, 2]. However, due to the strong nonlinearity and variability of market mechanisms, existing methods often struggle to capture complex causal structures accurately. Especially during extreme market conditions and unexpected events, traditional models based on static correlations suffer severe performance degradation. This exposes the urgent need for deeper causal modeling [3].

In recent years, the development of representation learning has provided new breakthroughs for financial data modeling. By automatically learning the latent feature representations, representation learning can effectively overcome the limitations of manual feature engineering in traditional factor discovery. However, current representation learning applications in finance mostly aim at improving prediction accuracy. They often lack in-depth exploration and utilization of the true causal relationships among variables. Since financial market data commonly include confounding factors, hidden variables, and spurious correlations, purely relying on traditional representation learning may easily lead to overfitting and poor interpretability. Therefore, incorporating causal inference into the representation learning process, to identify and model the causal structures among variables, is key to enhancing both the performance and robustness of return predictions.

Variational inference, as an efficient approximate reasoning method, demonstrates strong flexibility and scalability in modeling complex distributions. Combining variational methods with causal representation learning allows for maintaining learning efficiency while better revealing the latent causal mechanisms behind the data. By constructing a latent variable space and optimizing the structure and parameters of causal graphs through variational methods, it becomes possible to extract more interpretable and generalizable causal representations from high-dimensional, noisy financial data. This approach not only helps improve prediction accuracy but also significantly enhances model transferability and robustness in changing market environments. It provides a more reliable basis for financial decision-making [4, 5].

In market return prediction tasks, multi-factor models still dominate. Different factors reflect dynamic market changes from different perspectives. How to select key factors with real causal contributions from a wide range of possible influencing variables and then build interpretable and generalizable predictive models remains a major challenge. Methods based on variational causal representation learning can automatically discover hidden effective causal paths within the factor space. They can suppress the interference of redundant or false signals, thereby achieving causally driven multi-factor return modeling. This approach not only promises improved prediction performance but also provides strong technical support for practical applications, such as financial risk control and asset allocation strategy optimization.

In conclusion, integrating variational inference, representation learning, and causal reasoning into multi-factor market return prediction carries significant theoretical and practical value. On the one hand, it helps deepen the understanding of the intrinsic mechanisms of financial markets, promoting a shift from correlation-driven to causality-driven financial modeling. On the other hand, by building causal representations, it enables more stable and interpretable return predictions under high uncertainty, enhancing the practical applicability of models in real-world financial businesses. In the future, as financial data continue to grow in volume and complexity, exploring efficient and scalable causal representation learning methods will become a key development direction in financial artificial intelligence.

2. Related work

2.1 Variational Inference

Variational inference, as an important approximate inference method, shows excellent efficiency and flexibility in handling complex posterior distribution computations. When facing high-dimensional, nonlinear, and intractable posterior problems, variational inference transforms the inference task into an optimization problem[6]. It searches for a computable variational distribution to approximate the true posterior, effectively reducing computational complexity. The introduction of the Evidence Lower Bound (ELBO) allows the inference process to proceed efficiently using optimization techniques such as gradient descent. It ensures the quality of the approximation while greatly expanding the scale and complexity of applicable models [7]. Due to its good scalability and stability, variational inference has become a mainstream inference technique in areas such as probabilistic graphical models and deep generative models.

In financial time series modeling, traditional inference methods often struggle to deal with the high noise, strong nonlinearity, and complex latent structures common in real financial data [8]. This is because they rely heavily on strict assumptions about model structure and data distribution. Variational inference, by flexibly designing variational distributions, can better adapt to the latent structures and dynamic evolution patterns typical of financial data. It demonstrates higher robustness in tasks such as return prediction and risk assessment [9]. Furthermore, with the development of adaptive variational inference and black-box variational inference, variational methods can efficiently infer dynamic changes of latent factors in complex financial environments without relying heavily on detailed model information. This provides strong support for uncovering the underlying mechanisms of markets.

The combination of variational inference and deep learning has further advanced the innovative applications of representation learning and causal reasoning. By introducing neural networks to parameterize variational distributions within the variational framework, it becomes possible to flexibly model the latent causal structures of high-dimensional input [10]. At the same time, it maintains the differentiability and optimizability of the inference process. This combination not only enhances the expressive power of the representations but also provides new pathways for modeling complex causal relationships in financial data. In multi-factor market return prediction tasks, methods based on variational inference can effectively capture hidden causal associations between factors. This strengthens the model's ability to recognize heterogeneous market signals and lays a solid foundation for subsequent causal representation learning and predictive modeling.

2.2 Causal Inference

Causal inference, as a core method for understanding the mechanisms between variables, has been widely applied in many fields in recent years [11]. Unlike traditional correlation analysis, causal inference focuses on causal effects between variables. It studies the actual impact of one variable on another under external interventions or environmental changes. In the financial domain, market data are often filled with confounding factors and potential biases. Relying only on statistical correlation can easily lead to misjudgment and overfitting. Causal inference methods explicitly model the relationships between interventions and counterfactuals [12]. They can more accurately identify the key factors that truly drive market changes. This improves model interpretability and extrapolation ability and provides a stronger theoretical foundation for return prediction and risk management [13, 14].

In recent years, many causal inference methods that combine structured modeling and machine learning techniques have emerged. They have advanced the applications of causal discovery and causal effect estimation in complex data environments [15]. Especially in financial data that are high-dimensional, dynamic, and contain unobserved confounding variables, methods such as causal structure learning based on graphical models and counterfactual inference based on intervention modeling have shown unique advantages. Through causal inference, it is possible to effectively distinguish direct effects from indirect effects and eliminate spurious correlations. This provides feature representations with causal semantics for representation learning. The development of this direction not only enhances the robustness of predictive models in complex financial markets but also lays an important foundation for building intelligent financial systems with reasoning and decision-support capabilities.

2.3 Representation Learning

Representation learning aims to automatically extract useful feature representations from raw data so that downstream tasks can more effectively utilize the structural information within the data [16]. Compared to traditional methods relying on manual feature engineering, representation learning captures high-level semantics and hidden relationships through complex models such as deep networks. It has shown outstanding performance in tasks such as classification, prediction, and generative modeling [17, 18]. In financial market modeling, data are often noisy, high-dimensional, and nonlinear. Representation learning can automatically discover deep features that traditional financial factors fail to capture. This enhances the model's ability to perceive and adapt to market changes and provides stronger data-driven support for applications such as return prediction, risk control, and asset pricing.

However, most current mainstream representation learning methods are performance-driven and lack modeling of causal structures among features. They are easily disturbed by spurious correlations or environmental changes, which leads to performance degradation in new environments. Therefore, incorporating causal inference methods into the representation learning process to learn latent features with causal meaning has become an important research direction in recent years. Through causal representation learning, it is possible to extract features that are truly useful and robust for tasks. It also improves model interpretability and transferability. This is particularly suitable for environments like financial markets, which are highly dynamic and structurally complex. It provides a new modeling paradigm for tasks such as multi-factor return prediction.

3. Method

This study aims to model the deep causal relationship between market returns and multiple factors through variational causal representation learning. First, the overall architecture of the model is given, as shown in Figure 1.

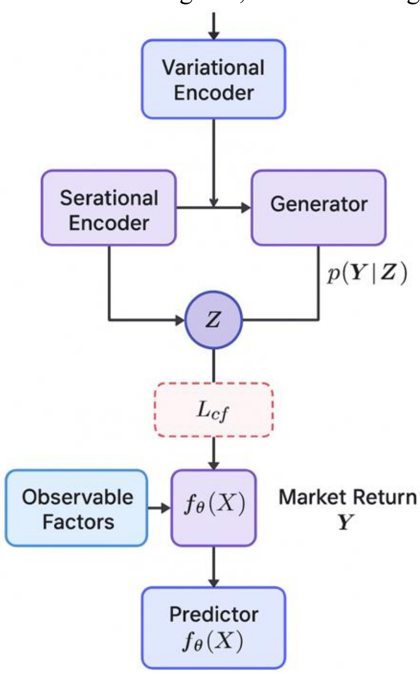

Figure 1. Overall model architecture diagram

The model architecture adopts a variational causal representation learning framework to encode observed factors into latent causal representations Z to capture hidden market generation mechanisms. By jointly training the variational encoder, generator, and predictor, the model not only achieves data reconstruction and causal consistency control, but also improves the ability to predict market returns. The counterfactual regularization term introduced in the architecture strengthens the causal interpretability of latent variables, making the prediction results more robust in a volatile market environment.

Assume that the observed data consists of explicit factor variables X, latent factors Z and market return Y, and there is the following generation process:

$$p(Y,X,Z) = p(Y \mid Z)p(X \mid Z)p(Z)$$

Where Z represents the implicit causal mechanism modeled by the latent variable. Since the true posterior distribution $p(Z \mid X,Y)$ is usually difficult to calculate directly, variational inference is introduced to approximate the true posterior using a parameterizable variational distribution $q_\phi(Z \mid X,Y)$.

The goal of variational inference is to maximize the evidence lower bound (ELBO) to minimize the Kullback-Leibler divergence between the variational distribution and the true posterior distribution. The specific optimization goal is:

$$L(\phi,\theta) = E_{q_\phi}[\log p_\theta(Y,X,Z) - \log q_\phi(Z \mid X,Y)]$$

In order to strengthen the causal interpretability of the latent variable Z, this method introduces a causal regularization term to encourage the learned representation to satisfy intervenibility and counterfactual consistency. Specifically, the counterfactual consistency loss is designed:

$$L_{cf} = E_{q_\phi}[d(f_\theta(do(X'=x')), f\theta(X'))]$$

Where $do(X'=x')$ represents external intervention on the input variables, $do(\cdot,\cdot)$ is the consistency measurement function, and f_θ is the yield prediction function. This loss ensures that the representation maintains reasonable changes under intervention conditions, which meets the basic requirements of causal reasoning.

The overall optimization objective integrates reconstruction loss, causal regularization term and prediction loss, which can be expressed as:

$$L_{total} = L(\phi,\theta) + \lambda_{cf} L_{cf} + \lambda_{pred} L_{pred}$$

Where $\lambda_{cf}, \lambda_{pred}$ is a hyperparameter that adjusts the contribution of each part, and the prediction loss L_{pred} is defined as the yield prediction error:

$$L_{pred} = E[(f_\theta(X) - Y)^2]$$

During the training process, the latent variable Z is sampled using the reparameterization technique to ensure gradient transferability. Finally, by jointly training the variational encoder, generator, and predictor, not only can the observed data be reconstructed and the underlying causal mechanism be captured, but the accuracy and robustness of market yield forecasts can also be effectively improved. The multi-factor modeling framework based on variational causal representation helps to screen out key factors with actual causal effects in high-dimensional financial data, suppress false signal interference, and provide a more explanatory and stable basis for financial decision-making.

4. Experimental Results

4.1 Dataset

This study uses the publicly available FRED-MD (Federal Reserve Economic Data - Monthly Database) as the main data source. The dataset is compiled by the Federal Reserve System of the United States. It covers macroeconomic time series from the 1960s to the present. It includes monthly economic variables such as employment, inflation, interest rates, money supply, consumption, production, and industrial indicators. It is widely used in economic forecasting and causal inference research.

The FRED-MD dataset is highly stable and representative. It contains around 100 macroeconomic variables. With its long-time span and rich set of indicators, it provides a comprehensive view of the dynamic behavior of the U.S. economy. These variables form a high-dimensional, dynamic, multi-factor feature space. This creates a complex and realistic experimental environment for causal representation learning. It helps uncover deep economic mechanisms that drive market behavior.

To meet model input requirements, all-time series were standardized, and missing values were filled during preprocessing. The time dimension was aligned to a unified monthly frequency. In the experiments, a fixed number of past periods of economic indicators were used as input variables. A representative economic measure, industrial production, was used as a proxy for returns. This setup was used to evaluate the model's capability in causal modeling and prediction.

4.2 Experimental Results

1) Experiments comparing this algorithm with other algorithms

In this section, this paper first gives the comparative experimental results of the proposed algorithm and other algorithms, as shown in Table 1.

Table 1: Comparative experimental results

Method	MSE	Causal Consistency	Robustness Score
Ours	0.017	0.91	0.85
TARNet[19]	0.025	0.79	0.70
CFR[20]	0.023	0.81	0.73
CEVAE[21]	0.024	0.82	0.72
DragonNet[22]	0.022	0.84	0.76

According to the experimental results, the model proposed in this study outperforms mainstream causal inference methods across all evaluation metrics. In terms of mean squared error (MSE), the proposed method achieves 0.017, which is significantly lower than other models. This indicates that the multi-factor prediction model based on variational causal representation learning has higher predictive accuracy in market return modeling. In contrast, traditional methods such as TARNet, CFR, CEVAE, and DragonNet show a certain degree of disadvantage in error control. This validates the importance of incorporating causal structure modeling to improve prediction precision.

For the causal consistency metric, the proposed method also achieves the highest score of 0.91, which is significantly higher than other models. This shows that the method can better capture stable causal relationships between factors and the target variable during modeling. It also avoids interference from spurious correlations in the prediction results. In comparison, although CEVAE and DragonNet have certain advantages in causal modeling, their consistency scores are still lower than that of the proposed method. This reflects that solely relying on latent variable inference or counterfactual modeling still has limitations.

In terms of robustness score, the proposed method reaches 0.85, far exceeding other baseline models. This further demonstrates the method's stability when facing input perturbations or market changes. Traditional models such as TARNet and CFR tend to exhibit fluctuations in predictive performance under different data disturbance conditions. In contrast, the proposed method, through causal regularization and latent space modeling, effectively enhances the model's extrapolation ability and noise resistance. It shows stronger generalization performance and practical application potential.

Overall, the experimental results fully demonstrate the effectiveness of integrating variational inference, causal reasoning, and deep representation learning. By properly constructing latent causal representations and reinforcing counterfactual consistency during training, the proposed method not only improves the accuracy of return prediction but also enhances model interpretability and adaptability in complex financial environments. These findings provide a new research direction and practical reference for optimizing financial prediction models based on causal mechanisms in the future.

2) Ablation experiment

In this section, this paper gives the impact of variational causal inference on model performance. The experimental results are shown in Table 2.

Table 2: Ablation experiment

Method	MSE	Causal Consistency	Robustness Score
Ours	0.017	0.91	0.85
w/o Causal Regularization	0.021	0.79	0.73
w/o Latent Variable Modeling	0.023	0.76	0.69
w/o Counterfactual Loss	0.020	0.81	0.74

The ablation study results show that each component plays an important role in improving the overall model performance. The complete model achieves the best results across all metrics. It reaches the lowest mean squared error (MSE) of 0.017, a causal consistency score of 0.91, and a robustness score of 0.85. This indicates that introducing causal regularization, latent variable modeling, and counterfactual loss has a significant positive impact on market return prediction tasks.

Specifically, after removing causal regularization, the MSE rises to 0.021, the causal consistency score drops to 0.79, and the robustness score falls to 0.73. This shows that the constraint from causal structures plays a key role in ensuring model generalization and prediction stability. The impact of removing latent variable modeling is even more pronounced. The MSE increases to 0.023, and both causal consistency and robustness scores decrease sharply. This reflects that latent causal representations are critical for capturing complex economic relationships.

In addition, after removing the counterfactual loss, the causal consistency score decreases but to a smaller extent, and the MSE also shows a slight increase. This suggests that counterfactual modeling provides auxiliary benefits for improving model stability and resisting spurious correlations. Overall, the ablation results demonstrate that latent variable modeling and causal regularization are the two core elements supporting the model's performance. Counterfactual consistency further enhances the model's robustness and interpretability.

3) Prediction performance analysis under different latent variable dimension settings

This paper also gives the prediction performance analysis under different latent variable dimension settings, and the experimental results are shown in Figure 2.

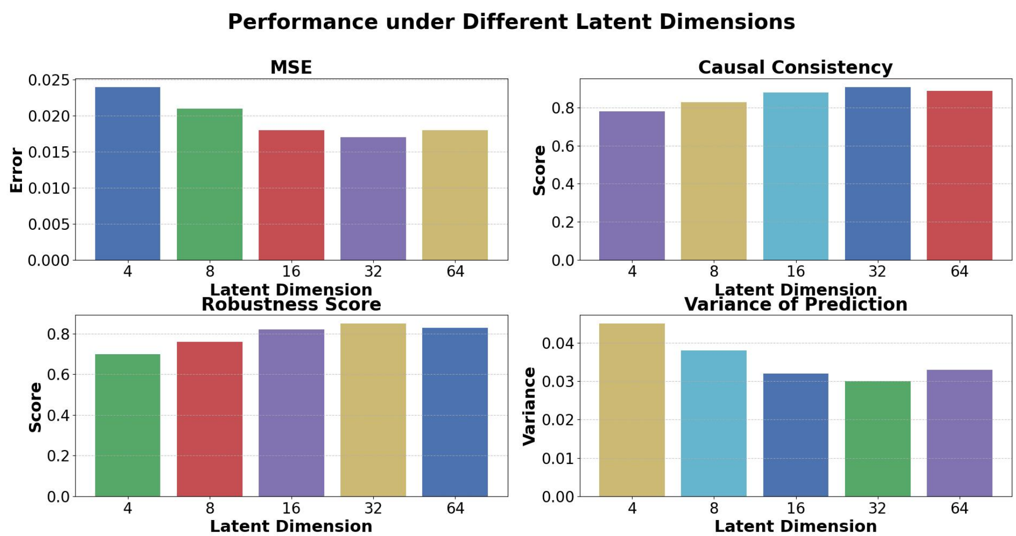

Figure 2. Performance Under Different Latent Dimensions

The results in the figure show that the dimension of latent variables has a significant impact on model performance. In terms of the MSE metric, as the dimension increases from 4 to 32, the error continuously decreases and reaches its best at 32 dimensions. This indicates that the model can more effectively express the causal structure of the data with a moderate dimension. A dimension that is too small limits the representation capacity, while a dimension that is too large, such as 64, may introduce redundant information and cause a slight increase in error.

For the Causal Consistency and Robustness Score metrics, the model shows a general trend of improvement as the latent variable dimension increases. Both metrics reach their peak at 32 dimensions. This suggests that higher dimensions help better capture stable causal relationships and improve the model's adaptability to environmental perturbations. However, at 64 dimensions, the scores slightly decline, indicating that excessive degrees of freedom may affect the controllability and generalization of the causal structure.

In addition, the prediction variance metric shows that as the latent variable dimension increases, the fluctuation of model outputs tends to decrease. This indicates that the model predictions become more stable at higher dimensions. Nevertheless, a slight increase in variance at 64 dimensions confirms the risk of overfitting or redundant interference when the dimension is too high. Overall, 32 is the optimal choice for the latent variable dimension in this experiment, achieving a good balance between accuracy, consistency, and stability.

4) Results of the cross-economic cycle forecasting capability verification experiment

Finally, this paper also presents the experimental results that verify the model's forecasting ability across different economic cycles. These results are summarized in Table 4. The evaluation demonstrates how well the proposed method adapts to varying market conditions, including expansion, recession, trough, and recovery phases, providing further evidence of the model's robustness and generalization capability in dynamic economic environments.

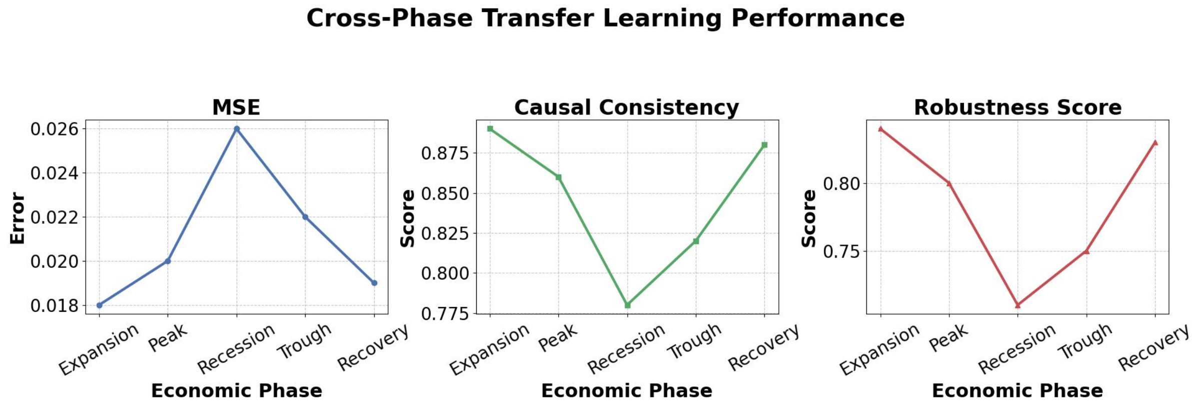

Figure 3. Results of the cross-economic cycle forecasting capability verification experiment

The experimental results show that the model's transfer performance varies significantly across different economic cycle stages. In the MSE plot, the model performs well during the "Expansion" and "Recovery" stages, with errors maintained at relatively low levels. However, during the "recession" stage, the error peaks. This indicates that the model faces greater prediction challenges during economic downturns, likely due to intensified shifts between the training distribution and the testing environment caused by dramatic market mechanism changes.

The causal consistency metric performs best during the "expansion" stage, with a score close to 0.88. This indicates that the model is able to capture stable and reliable causal structures when the economic environment is favorable and relatively stable. However, during the "recession" and "trough" stages, the causal consistency score drops significantly, reflecting the model's reduced ability to preserve accurate causal representations under heightened economic stress and volatility. This decline suggests that increased uncertainty and structural shifts in the market make it more difficult for the model to maintain consistent causal understanding. As the economy moves into the "recovery" stage, the causal consistency score rises again, showing that the model regains some of its causal inference capacity. This recovery demonstrates that the model has a certain level of adaptability in learning and maintaining causal relationships across different phases of the economic cycle.

The robustness score curve shows a similar trend. The model is most robust during the "Expansion" stage and experiences the sharpest decline during the "Recession" stage. This suggests that the model becomes more sensitive to input perturbations during periods of severe economic fluctuation. As the economy enters the "recovery" phase, the robustness score rises again. This reflects that the causally driven model has a certain ability to recover and generalize after cyclical shocks.

5. Conclusion

This study proposes a multi-factor market return prediction model based on variational causal representation learning. It integrates the strengths of variational inference, causal reasoning, and deep representation learning. The model effectively improves prediction accuracy and stability when dealing with high-dimensional financial data. By introducing latent causal variable modeling and counterfactual consistency regularization, the model can automatically extract factor representations with real economic significance. It also maintains good interpretability and causal consistency. This approach overcomes the adaptability limitations of traditional correlation-driven methods in non-stationary market environments. The experimental section validates the model's effectiveness from multiple perspectives. This includes comparisons with several classical causal inference models, ablation analyses of key modules, and evaluations of transfer performance across economic cycles. The results show that the proposed method achieves significant advantages in MSE, causal consistency, and robustness. In particular, it demonstrates stronger generalization ability under complex economic cycle changes. This performance improvement has important practical value for financial prediction tasks. It can provide more stable and reliable data support for applications such as investment strategy development, risk control, and macroeconomic policy management.

Moreover, the advantages of the model in causal mechanism modeling give it strong cross-task transfer potential. It is not only applicable to market return prediction but also generalizable to other economic variable modeling scenarios, such as credit risk assessment, policy impact analysis, and causal-driven intelligent decision systems. Its deep representation structure and causal consistency control mechanism lay a methodological foundation for building interpretable AI financial systems. This also promotes the practical application of causal learning in complex real-world systems. Future research can further extend the model's capabilities in heterogeneous data integration, cross-market modeling, and real-time online prediction. Introducing dynamic structural causal modeling and neural architecture search mechanisms may enhance the model's adaptability to sudden events and structural shifts. With the development of causal artificial intelligence, this study provides theoretical support and practical pathways for building more trustworthy, robust, and interpretable financial forecasting systems.

References

[1] Ican, Ozgur, and Taha Bugra Celik. "Stock market prediction performance of neural networks: A literature review." International Journal of Economics and Finance 9.11 (2017): 100-108.

[2] Zhang, Xincheng. "Country-level energy-related uncertainties and stock market returns: Insights from the US and China." Technological Forecasting and Social Change 204 (2024): 123437.

[3] Chopra, Ritika, Gagan Deep Sharma, and Vijay Pereira. "Identifying Bulls and bears? A bibliometric review of applying artificial intelligence innovations for stock market prediction." Technovation 135 (2024): 103067.

[4] Kelly, Bryan, Semyon Malamud, and Kangying Zhou. "The virtue of complexity in return prediction." The Journal of Finance 79.1 (2024): 459-503.

[5] Cakici, Nusret, et al. "Do anomalies really predict market returns? New data and new evidence." Review of Finance 28.1 (2024): 1-44.

[6] Bernardi, Mauro, Daniele Bianchi, and Nicolas Bianco. "Variational inference for large Bayesian vector autoregressions." Journal of Business & Economic Statistics 42.3 (2024): 1066-1082.

[7] Ober, Sebastian W. "Towards Improved Variational Inference for Deep Bayesian Models." arXiv preprint arXiv:2401.12418 (2024).

[8] Yin, Z., & Barucca, P. (2022). Variational heteroscedastic volatility model. arXiv preprint arXiv:2204.05806.

[9] Huang, Yating, Zhao Wang, and Cuiqing Jiang. "Diagnosis with incomplete multi-view data: A variational deep financial distress prediction method." Technological Forecasting and Social Change 201 (2024): 123269.

[10] Sahlström, T., & Tarvainen, T. (2023). Utilizing variational autoencoders in the Bayesian inverse problem of photoacoustic tomography. SIAM Journal on Imaging Sciences, 16(1), 89-110.

[11] Wang, Juanjuan, et al. "An ensemble model for stock index prediction based on media attention and emotional causal inference." Journal of Forecasting 43.6 (2024): 1998-2020.

[12] Xing, Wanqi, Chi Chen, and Lei Xue. "Deep Learning-Based Causal Inference Architecture and Algorithm between Stock Closing Price and Relevant Factors." Electronics 13.11 (2024): 2056.

[13] Bass, F. M., Jain, D., & Krishnan, T. (2000). Modeling the marketing-mix influence in new-product diffusion. New-product diffusion models, 99-122.

[14] Luo, Huishi, et al. "A survey on causal inference for recommendation." The Innovation 5.2 (2024).

[15] Naser, M. Z. "Causality and causal inference for engineers: Beyond correlation, regression, prediction and artificial intelligence." Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery 14.4 (2024): e1533.

[16] Khan, Muhib A., et al. "Evolving knowledge representation learning with the dynamic asymmetric embedding model." Evolving Systems 15.6 (2024): 2323-2338.

[17] Xu, X., Peng, H., & Chen, Y. (2021). Deep Switching State Space Model (DS $^ 3$ M) for Nonlinear Time Series Forecasting with Regime Switching. arXiv preprint arXiv:2106.02329.

[18] Liu, Jiexi, and Songcan Chen. "Timesurl: Self-supervised contrastive learning for universal time series representation learning." Proceedings of the AAAI conference on artificial intelligence. Vol. 38. No. 12. 2024.

[19] Chowdhury, Ranak Roy, et al. "Tarnet: Task-aware reconstruction for time-series transformer." Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data Mining. 2022.

[20] Kiriakidou, Niki, Ioannis E. Livieris, and Panagiotis Pintelas. "Mutual information-based neighbor selection method for causal effect estimation." Neural Computing and Applications 36.16 (2024): 9141-9155.

[21] Faghihi, Usef, Cyrus Kalantarpour, and Amir Saki. "Causal probabilistic based variational autoencoders capable of handling noisy inputs using fuzzy logic rules." Science and Information Conference. Cham: Springer International Publishing, 2022.

[22] Yin, M., Shi, C., Wang, Y., & Blei, D. M. (2024). Conformal sensitivity analysis for individual treatment effects. Journal of the American Statistical Association, 119(545), 122-135.