Article

Not peer-reviewed version

Empirical Evaluation of Quantum Kernel Learning for Sector Rotation Prediction: Using Fama–French 10 Industry Portfolios

Tohru Nakajima *

Posted Date: 10 February 2026

doi: 10.20944/preprints202602.0745.v1

Keywords: quantum kernel learning; sector rotation; financial time series; quantum machine learning; Fama–French industry portfolios

doi:10.20944/preprints202602.0745.v1

Article

Empirical Evaluation of Quantum Kernel Learning for Sector Rotation Prediction: Using Fama–French 10 Industry Portfolios

Tohru Nakajima

Graduate School of Agricultural and Life Sciences, The University of Tokyo;
nakajimatohru@g.ecc.u-tokyo.ac.jp

Abstract

Quantum Machine Learning (QML) is gaining attention for extracting nonlinear structures from high-dimensional data, yet its practical utility in financial forecasting remains insufficiently verified. This study empirically evaluates the effectiveness of quantum kernel learning for sector rotation prediction using daily returns from Fama–French 10 Industry Portfolios (2022–2024). We formulated a supervised classification task using a compact nine-feature set comprising mean return, volatility, win rate, excess return, MA5, MA20, max drawdown, 5-day return, and skewness. We compared classical Support Vector Classifiers (RBF and Linear kernels) against a fidelity quantum kernel (ZZFeatureMap) evaluated on AerSimulator using 1-nearest neighbor classification. Results indicate that while the classical RBF kernel achieved the highest average accuracy across all sectors (52.1%), the quantum kernel (50.0%) outperformed the classical linear kernel (48.4%) and achieved the highest accuracy in the Non-Durables (55.0%) and Manufacturing (52.2%) sectors, demonstrating superior stability. These findings suggest that the effectiveness of quantum kernels in the near term is not universal but sector-dependent and conditional. Although quantum methods currently trail strong classical baselines on average, future research should focus on advanced circuit design and hardware-aware evaluations to unlock their potential.

Keywords: quantum kernel learning; sector rotation; financial time series; quantum machine learning; Fama–French industry portfolios

1. Introduction

Quantum computing has garnered significant attention as a foundational technology capable of overcoming the computational complexity constraints of classical computing. In the field of machine learning, Quantum Machine Learning (QML) is anticipated to provide a new framework for extracting useful representations from data characterized by high dimensionality and non-linearity. Since the comprehensive outlook provided by Biamonte et al. [1], theoretical frameworks have advanced rapidly. Specifically, quantum kernel learning, which utilizes quantum circuits as feature maps to map classical data into a high-dimensional Hilbert space, has been established by Havlíček et al. [2] and Schuld et al. [3], demonstrating theoretical quantum advantage for data possessing specific algebraic structures.

However, in the current Noisy Intermediate-Scale Quantum (NISQ) era [4], the practical application of these algorithms faces significant challenges. As noted by Bharti et al. [5] and Cerezo et al. [6], issues such as qubit limitations, gate noise, and the "barren plateau" problem in variational circuits must be addressed. Furthermore, recent theoretical works present a more cautious view regarding quantum advantage. Liu et al. [7] demonstrated that merely utilizing quantum space is insufficient for learning speed advantage, while Huang et al. [8] emphasized the "power of data," revealing that the geometric structure of the kernel relative to the data volume is the dominant factor in learning performance. Consequently, there is a divergence between theoretical potential and practical utility, creating an urgent need to empirically verify under what specific conditions quantum kernels function effectively on realistic datasets [9,10].

Financial markets represent a prime testing ground for such empirical evaluations due to their inherent non-stationarity and noise [11]. While classical machine learning methods, such as deep neural networks and ensemble models, have become standard for asset pricing and forecasting [12–15], the field of "Quantum Finance" is exploring new possibilities for portfolio optimization and market prediction [16–19]. Although protocols for quantum forecasting exist [20,21], empirical studies that rigorously compare quantum kernels against strong classical baselines (e.g., optimized Support Vector Machines) using real-world market data remain limited.

This study focuses on "sector rotation," a widely used investment strategy based on the cyclical performance of industry sectors [22–24]. We aim to empirically evaluate the effectiveness of quantum kernel learning by applying it to the daily returns of the Fama–French 10 Industry Portfolios (2022–2024) [25]. By constructing a compact feature set and comparing a fidelity-based quantum kernel (ZZFeatureMap) with classical RBF and linear kernels, we assess whether quantum mapping offers superior decision boundaries in a noisy financial environment. Our principal conclusions indicate that while the classical RBF kernel achieves the highest average accuracy across all sectors, the quantum kernel demonstrates superior stability and outperforms classical methods in specific sectors such as Non-Durables and Manufacturing. These findings suggest that the advantage of quantum kernels in the NISQ era is not universal but rather sector-dependent, highlighting the need for task-specific circuit design.

2. Materials and Methods

2.1. Data Curation and Preprocessing

In this study, we utilized the "Fama–French 10 Industry Portfolios Daily Returns" provided by the Kenneth R. French Data Library [25]. The analysis period spans three years, from January 1, 2022, to December 31, 2024, comprising a total of 1,506 trading days. The ten target industries are NoDur (Consumer Non-Durables), Durbl (Consumer Durables), Manuf (Manufacturing), Enrgy (Energy), HiTec (High Technology), Telcm (Telecommunications), Shops (Wholesale and Retail), Hlth (Healthcare), Utils (Utilities), and Other.

To account for the non-stationarity of financial time series data and to strictly avoid look-ahead bias, the dataset was split while preserving chronological order. Specifically, the first 80% of the data (1,152 days) was designated as the training dataset, and the subsequent 20% (289 days) was allocated as the testing dataset. Regarding feature normalization, Z-score standardization (mean of 0, variance of 1) was performed using statistics (mean and standard deviation) derived exclusively from the training data. The same transformation parameters were then applied to the test data to prevent data leakage.

2.2. Feature Engineering

We constructed nine distinct feature variables based on the daily return series of the past 60 trading days for each industry to capture short-term momentum, distribution shape, downside risk, and trend components. All features were Z-score normalized using the statistics of the training dataset before being used as input for the kernel models. The definitions of the features are as follows:

- **Mean Return:** The arithmetic mean of daily returns over the past 60 days, representing the average performance level of the industry during the period.
- **Volatility:** Calculated as the unbiased standard deviation of daily returns over the past 60 days. This indicator quantifies the dispersion of returns and the magnitude of risk during the period.
- **Win Rate:** Defined as the proportion of trading days recording a positive return within the past 60 days. This measures the consistency of performance rather than the magnitude of returns.

- **Excess Return:** The average value over the past 60 days of the difference between the industry's daily return and the market average return (calculated as the equal-weighted average of all 10 industries). This evaluates performance relative to the overall market.
- **MA5 (5-Day Moving Average):** The simple moving average of daily returns over the most recent 5 days. This corresponds to one trading week and is used to smooth and capture extremely short-term trend components.
- **MA20 (20-Day Moving Average):** The simple moving average of daily returns over the most recent 20 days. This corresponds to approximately one trading month and serves as an indicator to evaluate medium-term market trends by removing short-term noise over a span longer than MA5.
- **Max Drawdown:** An indicator of downside risk defined as the maximum percentage decline from peak to trough in cumulative returns (calculated via compounding) over the past 60 days. In implementation, the cumprod() function was used to calculate the maximum erosion rate of asset value within the period.
- **5-Day Return:** The sum of returns over the most recent 5 days, reflecting short-term momentum. While MA5 observes the average level, this metric was adopted to directly represent the cumulative magnitude of fluctuation.
- **Skewness:** The skewness (third moment) of the return distribution over the past 60 days. This represents the deviation from a normal distribution; specifically, negative skewness suggests a distribution shape with a higher risk of crashes (sudden drops) that deviate significantly from the mean.

2.3. Label Generation

The prediction task in this study is formulated as a binary classification problem to determine whether the returns of each sector will outperform the market average in the short term. The objective is to remove market-wide fluctuations (systematic risk) and capture relative performance attributable to sector-specific factors (idiosyncratic risk).

The prediction horizon is defined as the five-day period starting from the next business day after time t (from $t+1$ to $t+5$). First, the cumulative return $R\_i,t(5)$ for sector i over this period is calculated as the simple sum of daily returns r_i,τ using the following equation:

$$R_{(i,t)(5)} = \sum_{\tau=1}^{5} r_{i,t+\tau}$$

Similarly, for the benchmark market average return, we adopted the equal-weighted average of the Fama–French 10 Industry Portfolios. The daily market return r_M,τ at time τ is defined as:

$$r_{M,\tau} = \frac{1}{10}\sum_{j=1}^{10} r_{j,\tau}$$

Using this, the cumulative return for the overall market $R_{M,t}(5)$ over the same period is calculated in the same manner:

$$R_{M,t}(5) = \sum_{\tau=1}^{5} r_{M,t+\tau}$$

The final ground-truth label $y_{i,t} \in \{0,1\}$ is determined based on whether the sector's cumulative return exceeds the market's cumulative return:

$$y_{i,t} = \begin{cases} 1 & \text{if } R_{i,t}(5) > R_{M,t}(5) \\ 0 & \text{otherwise} \end{cases}$$

Here, $y_{i,t} = 1$ signifies that the sector exhibits performance superior to the market average ("Outperform"), while 0 signifies performance equal to or below the market average ("Underperform"). Note that for the calculation of period returns, we utilized arithmetic sums (simple interest) rather than strict geometric means (compound interest). This decision was made because the target period is short (5 trading days) and daily return values are sufficiently small, such that the linear approximation $\ln(1+r) \approx r$ holds. Thus, the arithmetic sum was judged to be appropriate given the trade-off between computational cost and precision.

2.4. Data Splitting and Normalization

Data splitting was performed chronologically: the first 80% was designated as the training data, and the latter 20% as the test data (random shuffling was not performed). For feature normalization, a StandardScaler was fitted on the training data, and the same transformation was applied to the test data. Using the mean $¥mu$ and standard deviation $¥sigma$ of the features during the training period, each feature vector x is standardized as $x' = \frac{x-\mu}{\sigma}$ (element-wise division). This ensures that scaling statistics are estimated based solely on information from the training period, thereby avoiding the leakage of future information.

2.5. Quantum Kernel Learning (FidelityQuantumKernel + ZZFeatureMap)

In the quantum kernel method, an input vector x is mapped to a quantum state via a quantum circuit $U(x)$, defined as $|\varphi(x)\rangle = U(x)|0\rangle^{\otimes n}$. The kernel value between two points x and x' is defined as the fidelity between the corresponding quantum states: $k(x, x') = |\langle\varphi(x)|\varphi(x')\rangle|^2$ [2]. This introduces a non-linear similarity measure based on the inner product in the quantum state space, enabling learning using a kernel matrix in the same manner as classical kernel methods.

This implementation was constructed using the Qiskit framework, divided into two phases: Quantum Processing and Classical Post-processing. First, in the Quantum Processing phase, we adopted the ZZFeatureMap for the feature map. The number of repetitions, reps, is a parameter directly linked to circuit depth and the computational cost of fidelity estimation. This choice of circuit configuration is closely tied to the design of our feature space. We intentionally utilized a compact set of hand-crafted features to align with the constraints of current Noisy Intermediate-Scale Quantum (NISQ) devices. High-dimensional feature maps require deeper circuits with more entangling gates, which significantly increases susceptibility to noise and decoherence. Therefore, to balance the expressiveness of the quantum state space with the necessity of maintaining circuit fidelity, reps was fixed at 1.

In the subsequent Classical Post-processing phase, the kernel matrix was calculated based on the fidelity between states obtained from the quantum circuit. We employed AerSimulator as the backend and calculated fidelity via statevector simulation using the ComputeUncompute algorithm, which is the default implementation of FidelityQuantumKernel.

For the final classification, we adopted the 1-nearest neighbor (1-NN) algorithm. The reason for intentionally selecting 1-NN over learners involving optimization, such as Support Vector Machines (SVM), was to exclude the influence of additional optimization and to directly evaluate the effect of the geometric structure (similarity) of the feature space induced by the quantum kernel itself. While practical financial applications ultimately require assessment via economic metrics such as risk-adjusted returns and transaction costs, these metrics introduce external variables independent of the kernel's intrinsic performance. Consequently, this study prioritizes classification accuracy to validate the fundamental separability of the data mapped by the quantum kernel as a prerequisite for future economic backtesting. The specific prediction procedure is as follows: for a test sample x, the kernel values with the training reference set $\{x_j\}$ are calculated, and the label y_{j*} of the training data point with the highest similarity (maximum kernel value) is assigned as the predicted label:

$$\hat{y}(x) = y_{j^*}, \quad j^* = \arg\max_j k(x, x_j)$$

It should be noted that the computational cost of the quantum kernel tends to increase quadratically with the size of the reference set. Considering time constraints, we used a subset of 100

doi:10.20944/preprints202602.0745.v1

samples extracted from the training data as the reference set for the quantum kernel calculation, rather than the full dataset (whereas the full dataset was used for classical kernel calculations). To ensure reproducibility, these 100 samples were not selected via random sampling but were deterministically selected as the first 100 samples (indices 0–99) from the data sorted chronologically within the training interval.

2.6. Classical Kernels (SVC: RBF / Linear)

We adopted the Support Vector Classifier (SVC) from scikit-learn as the classical method for comparison. The RBF kernel enables non-linear separation via Gaussian kernels and serves as a strong baseline in financial non-linear pattern recognition. Conversely, the linear kernel provides the simplest decision boundary and is useful as a control to measure the contribution of non-linearity. Both models were trained using the full training dataset (approximately 1,152 samples). Hyperparameters were based on implementation defaults to avoid excessive tuning, prioritizing the reproducibility and interpretability of the comparative experiment.

2.7. Evaluation Metrics and Implementation

We adopted Accuracy as the metric for model performance evaluation. In addition to the prediction accuracy for each sector, we calculated the average accuracy and standard deviation across all 10 sectors to evaluate the model's generalization performance and stability.

$$\boldsymbol{Accuracy} = \left(\frac{\mathbf{1}}{\boldsymbol{N}}\right)\sum_{\boldsymbol{n}=\mathbf{1}}^{\boldsymbol{N}} \boldsymbol{I}(\boldsymbol{y_n} = \hat{\mathrm{y}}_{\boldsymbol{n}})$$

Here, $\boldsymbol{I}(\cdot)$ represents the indicator function.

All experiments in this study were conducted in a Python 3.12 environment. For quantum kernel calculations, we used Qiskit Machine Learning (v0.8.4) with AerSimulator as the backend, and for the implementation of classical classifiers (SVC), we used scikit-learn (v1.7.2). To ensure experimental reproducibility, data splitting and the extraction of the 100 samples for quantum learning were performed based on deterministic rules following chronological order, and a fixed seed was set for any processes requiring random numbers.

3. Results

3.1. Sector-Specific Classification Accuracy

This section details the experimental results obtained using the Fama–French 10 Industry dataset. We compared three methods: the classical RBF kernel (Classical RBF), the classical linear kernel (Classical Linear), and the quantum kernel (Quantum ZZ). Classification accuracy on the test data was employed as the evaluation metric.

Regarding the classification accuracy for each of the 10 Fama–French industries, the quantum kernel (Quantum ZZ) demonstrated higher performance than classical methods in two sectors: NoDur (Consumer Non-Durables) and Manuf (Manufacturing). Specifically, for NoDur, Quantum ZZ recorded 55.0%, surpassing Classical Linear (54.0%) and Classical RBF (52.9%). Furthermore, in Manuf, Quantum ZZ achieved 52.2%, whereas Classical RBF and Classical Linear were limited to 47.8% and 45.7%, respectively. As indicated by these results, the improvement by the quantum kernel reached a maximum of 6.5 percentage points in the manufacturing sector, confirming a significant advantage.

Conversely, classical methods exhibited superiority in other sectors. Notably, in HiTec (High Technology), Classical Linear recorded the highest accuracy of 59.9% across all experiments. This result suggests that a linearly separable structure dominates over complex non-linear mappings in this sector. Additionally, for Enrgy (Energy), Classical RBF recorded the highest value at 58.1%. Thus,

doi:10.20944/preprints202602.0745.v1

it was confirmed that the optimal kernel method is not uniform but depends strongly on sector-specific data characteristics, particularly linear separability and volatility properties.

3.2. Overall Average Accuracy and Stability

Next, we calculated the mean accuracy and standard deviation to evaluate the generalization performance and stability of the models across all 10 industries. In terms of mean accuracy, Classical RBF performed best at 52.1%, followed by Quantum ZZ at 50.0% and Classical Linear at 48.4%. Although the mean accuracy of the quantum kernel is 2.1 percentage points lower than that of the classical RBF, it remains 1.6 percentage points higher than that of the linear kernel.

Of particular note is the difference in stability represented by the standard deviation. The standard deviation of Quantum ZZ remained at 3.2%, the smallest among the three methods. In contrast, the standard deviation for Classical Linear was notably large at 7.0%, indicating that performance fluctuates significantly depending on the industry. The standard deviation for Classical RBF was 4.0%. These results indicate that the quantum kernel is less prone to extreme performance degradation due to differences in data distribution across sectors, possessing relatively stable learning characteristics.

4. Discussion

In this study, we formulated sector rotation prediction as a binary classification problem to determine whether each industry would outperform the market average over the subsequent five trading days. We compared classical kernels (RBF and Linear) with a quantum kernel (fidelity kernel based on the ZZFeatureMap). The results indicated that the classical RBF kernel achieved the highest average classification accuracy across all 10 industries at 52.1%, followed by the quantum ZZ kernel at 50.0% and the classical linear kernel at 48.4%. Consequently, under the current experimental settings, we cannot conclude that the quantum kernel universally outperforms the strong classical RBF baseline in terms of average performance.

However, a notable finding is that the inter-sector accuracy variance (standard deviation) of the quantum ZZ kernel was the lowest at 3.2%, indicating that performance fluctuations across industries were suppressed compared to the classical RBF (4.0%) and classical linear (7.0%) kernels. Financial time series data are characterized by strong non-stationarity, and the impact of exogenous shocks and regime changes varies by industry. Therefore, not only average accuracy but also "stability that resists extreme performance degradation" holds significant practical importance [11]. Although this study focuses on classification accuracy as a primary metric for signal detection, this observed stability—low performance variance—suggests that quantum kernels may offer economic value in the context of risk management by potentially reducing portfolio drawdown risk.

Regarding sector-specific performance, the quantum ZZ kernel achieved the highest accuracy in two industries: Consumer Non-Durables (NoDur) and Manufacturing (Manuf). In particular, an improvement of 4.4 percentage points over the classical RBF was observed in the Manufacturing sector. It is possible that in these industries, short-term relative strength manifests more consistently as a deviation from the market average, suggesting that the feature engineering based on statistics from the past 60 trading days and short-term momentum functioned effectively.

Nevertheless, due to computational cost constraints, the quantum approach in this study was limited to a learning reference set of 100 samples and utilized a 1-nearest neighbor (1-NN) classifier. Consequently, the performance of the quantum ZZ kernel is simultaneously influenced by (i) the representativeness of the reference set, (ii) the nature of 1-NN which relies on local neighborhoods, and (iii) the expressibility constraints of a shallow quantum circuit (reps=1). It is important to note that we intentionally employed this compact feature set and shallow circuit depth to align with the hardware limitations of the NISQ era. While high-dimensional feature spaces might offer greater expressiveness, they require deeper circuits that introduce significant noise and decoherence errors. Therefore, we prioritized the mitigation of noise over high-dimensional expressiveness in this initial evaluation. While these constraints may work disadvantageously for industries with complex distributions and high time-variability, the advantages of the quantum feature map likely surfaced in industries where short-term relative structures are stable. Thus, it is appropriate to interpret the results of this study not as "uniform quantum advantage" but as "conditional (sector-dependent) advantage."

There are several clear limitations to this study. First, the classifiers and training conditions were not perfectly aligned between the classical and quantum methods. While a Support Vector Classifier was used for the classical approach, the quantum approach adopted 1-nearest neighbor to directly evaluate the similarity structure defined by the quantum feature map; strictly speaking, this does not constitute an end-to-end performance comparison. Second, the design of the quantum kernel was limited to the ZZFeatureMap (reps=1), and optimization of circuit depth or entangling structures was not performed. Third, the learning reference set for the quantum method was limited to 100 samples due to computational time constraints.

Future research directions should include incorporating the quantum kernel as a precomputed kernel into a Support Vector Classifier to use the same classifier structure for both classical and quantum methods, thereby eliminating the influence of classifier discrepancies. Additionally, improving scalability through the gradual expansion of the reference set size and the introduction of kernel approximation methods, such as the Nyström approximation or advanced sampling strategies, is a key challenge. Furthermore, while classification accuracy serves as a fundamental validation of the kernel's pattern recognition capability, bridging the gap to practical investment strategies is essential. Future work must extend the evaluation to economic value by connecting classification results to actual sector rotation strategies and verifying practical feasibility using metrics such as cumulative returns, Sharpe ratios, and maximum drawdowns incorporating transaction costs.

5. Conclusions

In this study, to verify whether quantum kernel learning can demonstrate practical advantages for sector rotation prediction in financial data, we conducted a comparative evaluation of quantum and classical kernel methods using daily returns from the Fama–French 10 Industry Portfolios. We specifically designed a compact feature set suitable for NISQ devices to evaluate the fundamental geometric properties of the quantum kernel under noise-constrained conditions. The primary findings are as follows. First, the quantum kernel achieved the highest accuracy in 2 out of 10 industries (Consumer Non-Durables and Manufacturing), showing a significant improvement of +4.4 percentage points over the classical RBF kernel specifically in the Manufacturing sector. This suggests that the similarity structure defined by the quantum feature map can function effectively for identifying short-term relative strength in certain industries. Second, regarding the average accuracy across all industries, the classical RBF kernel performed best, and the quantum kernel did not surpass it. Third, the quantum kernel exhibited the smallest performance variance among industries, demonstrating relatively stable behavior across sectors.

These results indicate that the effectiveness of quantum kernel learning should be understood not as a "universal quantum advantage" that appears uniformly in overall averages, but as a conditional advantage that manifests under specific industries or market conditions. Therefore, from the perspective of practical application, it is appropriate to position the quantum kernel not as a universal replacement for classical methods, but as a complementary tool to be deployed selectively according to industry characteristics and market phases. Future work will focus on scaling the feature dimensions and reference sets while addressing noise mitigation, as well as evaluating the economic utility of the model through rigorous backtesting that accounts for market friction.

Author Contributions: Conceptualization, T.N.; methodology, T.N.; software, T.N.; validation, T.N.; formal analysis, T.N.; investigation, T.N.; resources, T.N.; data curation, T.N.; writing—original draft preparation, T.N.; writing—review and editing, T.N.; visualization, T.N.; supervision, T.N.; project administration, T.N.; funding acquisition, T.N.

doi:10.20944/preprints202602.0745.v1

The author has read and agreed to the published version of the manuscript.

Funding: This research was funded by the Foundation for the Promotion of Science and Technology (FOST) Grant Number 20241217 and the Yanmar Resource Circulation Support Organization.

Institutional Review Board Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Acknowledgments: The author acknowledges the support provided by the Foundation for the Promotion of Science and Technology (FOST) and the Yanmar Resource Circulation Support Organization.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Biamonte, J.; Wittek, P.; Pancotti, N.; Rebentrost, P.; Wiebe, N.; Lloyd, S. Quantum machine learning. *Nature* **2017**, *549*, 195–202, doi:10.1038/nature23474.
2. Havlíček, V.; Córcoles, A.D.; Temme, K.; Harrow, A.W.; Kandala, A.; Chow, J.M.; Gambetta, J.M. Supervised learning with quantum-enhanced feature spaces. *Nature* **2019**, *567*, 209–212, doi:10.1038/s41586-019-0980-2.
3. Schuld, M.; Killoran, N. Quantum machine learning in feature Hilbert spaces. *Phys. Rev. Lett.* **2019**, *122*, 040504, doi:10.1103/PhysRevLett.122.040504.
4. Preskill, J. Quantum Computing in the NISQ era and beyond. *Quantum* **2018**, *2*, 79, doi:10.22331/q-2018-08-06-79.
5. Bharti, K.; Cervera-Lierta, A.; Kyaw, T.H.; Haug, T.; Alperin-Lea, S.; Anand, A.; Degroote, M.; Heimonen, H.; Kottmann, J.S.; Menke, T.; et al. Noisy intermediate-scale quantum algorithms. *Rev. Mod. Phys.* **2022**, *94*, 015004, doi:10.1103/RevModPhys.94.015004.
6. Cerezo, M.; Arrasmith, A.; Babbush, R.; Benjamin, S.C.; Endo, S.; Fujii, K.; McClean, J.R.; Mitarai, K.; Yuan, X.; Cincio, L.; et al. Variational quantum algorithms. *Nat. Rev. Phys.* **2021**, *3*, 625–644, doi:10.1038/s41567-021-01250-3.
7. Liu, Y.; Arunachalam, S.; Temme, K. A rigorous and robust quantum speed-up in supervised machine learning. *Nat. Phys.* **2021**, *17*, 1013–1017, doi:10.1038/s41567-021-01287-z.
8. Huang, H.-Y.; Kueng, R.; Preskill, J. The power of data in quantum machine learning. *Nat. Commun.* **2021**, *12*, 2631, doi:10.1038/s41467-021-22563-5.
9. Schuld, M. Supervised quantum machine learning models are kernel methods. *arXiv* **2021**, arXiv:2101.11020.
10. Gentinetta, G.; Thomsen, A.; Eisert, J.; Briegel, H.J. The complexity of quantum support vector machines. *arXiv* **2022**, arXiv:2207.03058.
11. Ang, A.; Timmermann, A. Regime changes and financial markets. *Annu. Rev. Financ. Econ.* **2012**, *4*, 313–337, doi:10.1146/annurev-financial-110311-101750.
12. Dixon, M.; Klabjan, D.; Bang, J.H. Classification-based financial markets prediction using deep neural networks. *Algorithmic Finance* **2017**, *6*, 67–77, doi:10.3233/AF-170196.
13. Gu, S.; Kelly, B.; Xiu, D. Empirical asset pricing via machine learning. *Rev. Financ. Stud.* **2020**, *33*, 2223–2273, doi:10.1093/rfs/hhaa009.
14. Fischer, T.; Krauss, C. Deep learning with long short-term memory networks for financial market predictions. *Eur. J. Oper. Res.* **2018**, *270*, 654–669, doi:10.1016/j.ejor.2017.12.043.
15. Sezer, O.B.; Gudelek, M.U.; Ozbayoglu, A.M. Financial time series forecasting with deep learning: A systematic literature review: 2005–2019. *Appl. Soft Comput.* **2020**, *90*, 106181, doi:10.1016/j.asoc.2020.106181.
16. Orús, R.; Mugel, S.; Lizaso, E. Quantum computing for finance: Overview and prospects. *Rev. Phys.* **2019**, *4*, 100028, doi:10.1016/j.revip.2019.100028.
17. Bouland, A.; van Dam, W.; Joorati, H.; Kerenidis, I.; Prakash, A. Prospects and challenges of quantum finance. *arXiv* **2020**, arXiv:2011.06492.

doi:10.20944/preprints202602.0745.v1

18. Egger, D.J.; Gambella, C.; Mareček, J.; McFaddin, S.; Mevissen, M.; Raymond, R.; Simonetto, A.; Woerner, S.; Yndurain, E. Quantum computing for finance: State-of-the-art and future prospects. *IEEE Trans. Quantum Eng.* **2020**, *1*, 3101724, doi:10.1109/TQE.2020.3030314.
19. Herman, D.; Googin, C.; Liu, X.; Galda, A.; Safro, I.; Sun, Y.; Pistoia, M.; Céleri, Y. A Survey of Quantum Computing for Finance. *arXiv* **2022**, arXiv:2201.02773.
20. Kyriienko, O.; Magnusson, E.B.; Diswal, A. Protocols for programmable quantum machine learning and quantum forecasting. *Phys. Rev. A* **2021**, *103*, 052416, doi:10.1103/PhysRevA.103.052416.
21. Mugel, S.; Kuchkovsky, C.; Sanchez, E.; Simeon, S.; Suzuki, Y.; Lizaso, E. Dynamic portfolio optimization with real datasets using quantum processors and quantum-inspired tensor networks. *Phys. Rev. Res.* **2022**, *4*, 013006, doi:10.1103/PhysRevResearch.4.013006.
22. Conover, C.M.; Jensen, G.R.; Johnson, R.R.; Mercer, J.M. Sector rotation and monetary conditions. *J. Invest.* **2008**, *17*, 34–46, doi:10.3905/joi.2008.17.1.34.
23. Moskowitz, T.J.; Grinblatt, M. Do industries explain momentum? *J. Finance* **1999**, *54*, 1249–1290, doi:10.1111/0022-1082.00146.
24. Stangl, T.; Jacobsen, B.; Visaltanachoti, N. Sector rotation over business cycles. *Financ. Anal. J.* **2018**, *74*, 55–72, doi:10.2469/faj.v74.n1.3.
25. Fama, E.F.; French, K.R. Common risk factors in the returns on stocks and bonds. *J. Finance* **1993**, *48*, 3–56, doi:10.1111/j.1540-6261.1993.tb04702.x.