Quantum-enhanced versus classical Support Vector Machine: An application to stock index forecasting

Mathieu Mercadier*

DCU Business School

November 13, 2023

Draft. Not edited.

Please do not circulate.

Abstract

This paper synthesizes the quantum machine learning literature, focusing on quantum Support Vector Machine (SVM). Although its efficiency is recognized theoretically, it has limitations in practice, for instance, it is not mature enough for financial applications. Therefore, empirically, this study provides two experiments in which the quantum-enhanced SVM, using the quantum kernel estimator, is compared with the classical SVM. According to standard performance metrics, the current quantum-enhanced SVM does not show superior performances in forecasting the movement direction of stock market indexes. To the best of my knowledge, this study is a pioneer attempt applying this quantum algorithm in stock market index forecasting which provides insight to financial researchers and practitioners.

Keywords: Quantum Support Vector Machine; Support Vector Machine; Stock index forecasting; Forecasting practice; Financial markets

JEL Classification: C38; C45; G17

*Email: mathieu.mercadier@dcu.ie; DCU Business School, Dublin City University, DUBLIN 9, IRELAND.

1 Introduction

Machine learning aims at predicting and identifying patterns from data sets within a reasonable time. Data sets' size and dimension gradually increase over time, inexorably leading to calculation slowdown. An efficient solution to circumvent this issue potentially arises from an emerging technology: quantum computing. Specifically, quantum computers are expected to surpass the computational capabilities of classical computers during this decade (Pistoia et al., 2021).

Along with universities and research institutions, in the industry, there are several prominent companies in the rivalry of quantum computing development, such as IBM, Google, Alibaba, Microsoft, D-Wave, Rigetti Computing, IonQ, Xanadu, Fujitsu, etc., and, more recently, Honeywell (Arute et al., 2019; Zhong et al., 2021; Zhu et al., 2022b; Madsen et al., 2022). Google reaching quantum supremacy (Arute et al., 2019) in the sense of Preskill (2012) is criticized by Pednault et al. (2019) and Zhou et al. (2020). More recently, the enriching debate between classical and quantum supremacy continues with dynamic and intense exchanges (Kim et al., 2023; Tindall et al., 2023; Begušió and Chan, 2023).

The governments also show a growing interest in quantum technologies. France announced in January 2021 an investment plan of 1.8 billion euros over five years[1]. More recently in August 2022, the United States signed the CHIPS and Science Act 2022 that authorizes new investments in core quantum research programs[2]. And China is by far the biggest investor in quantum computing with USD 15 billions[3]. These investments play the role of incentives that promote the development of quantum computing, for instance, startup companies such as Alice&Bob compete in raising venture capital financing[4]. However, the investors are cautious in the rapid growth stage of industry lifecycle and await the shakeout stage where uncompetitive companies are eliminated in the rivalry, which puts pressure on quantum computing companies. For instance, some unicorn startups, namely Rigetti Computing, IonQ and D-Wave witnessed significant market capitalization drop by early 2023. Facing these pressures, Rigetti Computing is at the risk of being delisted[5]. Additionally, quantum supremacy threatens the security of the financial system. Consequently, JPMorgan Chase hired quantum cryptography scientists to

[1] https://news.cnrs.fr/articles/french-research-at-the-heart-of-the-quantum-plan
[2] https://www.quantum.gov/quantum-in-the-chips-and-science-act-of-2022
[3] https://www.mckinsey.com/featured-insights/sustainable-inclusive-growth/chart-of-the-day/betting-big-on-quantum
[4] https://pitchbook.com/newsletter/quantum-computing-startup-alicebob-raises-27m
[5] https://www.hpcwire.com/2023/02/03/quantum-computing-firm-rigetti-faces-delisting

develop quantum-resistant communication networks[6]. In July 2023, HSBC joined a quantum-secure network anticipating quantum-enabled cyber-threats[7].

Pistoia et al. (2021) state that finance is projected to be the first industry sector to reap the advantages of quantum computing for effective solutions. According to Nurdin (2019), randomness is implicit at the atomic scale and the measurement of a quantum mechanical system yields a random result. Financial studies shed the spotlight on capitalizing on the inherent stochastic properties of quantum mechanics (Montanaro, 2015; Rebentrost et al., 2018; Orùs et al., 2019; Alcazar et al., 2020). In addition, academics are motivated by another property which is the exponential speed-ups enabled by quantum machine learning (Lloyd et al., 2013; Rebentrost et al., 2014; Li et al., 2015; Biamonte et al., 2017; Rebentrost and Lloyd, 2018; Egger et al., 2020).

But in practice, beyond the need for financial practitioners' training, the technology itself is not yet mature enough for wide-scale industrial usage. Quantum computing is currently in the "noisy intermediate-scale quantum" (NISQ) era (Schuld and Petruccione, 2018; Park et al., 2020b). Access to actual NISQ devices is granted by IBM via its cloud solution. Thus, as in many recent papers, this study conducts financial experiments through the IBM Quantum Experience (Hebenstreit et al., 2017; Harper and Flammia, 2019; Martin et al., 2021; Mugel et al., 2022; Wilkens and Moorhouse, 2023).

This paper proposes two experiments forecasting the movement direction of stock market indexes using a supervised learning algorithm. The algorithm chosen for these experiments is the Support Vector Machine (SVM). This algorithm, robust to overfitting, can be used for classification tasks, effectively solve linear and non-linear problems, and perform well even with small data sets. Moreover, it finds extensive application in stock prediction (Gong et al., 2016; Chen and Hao, 2017; Ismail et al., 2020; Akyildirim et al., 2022; Kang et al., 2023), and recent research in forecasting (Richardson et al., 2021; Bas et al., 2021; Zhu et al., 2022a, 2023; Luo et al., 2023) and operations (Ben-Tal et al., 2015; Maldonado et al., 2017; Jiménez-Cordero et al., 2021; Benítez-Peña et al., 2023) proposes customizations and applications of this algorithm across various contexts. Besides, this study deals with a "quantum-enhanced" version of the SVM in which the quantum device handles the application of the kernel function to the data. In this case, the kernel function is referred as "quantum kernel estimator" (QKE).

[6]Cf. https://cde.nus.edu.sg/news-detail/cde-quantum-security-expert-charles-lim-joins-jpmorgan
[7]https://www.theregister.com/2023/07/06/hsbc_vodafone_quantum_security

QKE has an interesting property in not requiring any γ hyper-parameter setting, compared to the commonly used non-linear kernels. Indeed, the smaller the number of required parameters, the smaller the risk of incorrect settings. Additionally, the quantum-enhanced SVM does not require a choice among various non-linear kernels contrary to the SVM.

Besides the sensitivity to noise of quantum computers, disadvantages of the cloud-based technology lie in the restricted number of freely publicly usable qubits and long waiting queues. Therefore, the empirical experiments cannot be built on large samples, and they are limited to binary classifications. With regard to these limitations, the study attempts a pilot work to assess the level of performance that financial practitioners can expect using quantum-enhanced SVM. To do so, I compare the current predictive power of the quantum-enhanced SVM, to the predictive power of the SVM in practical financial contexts. So far, few empirical experiments used quantum-enhanced SVM but mainly on separable data sets (Li et al., 2015; Havenstein et al., 2018; Sarma et al., 2019; McRae and Hilke, 2020; Peters et al., 2021). In the first experiment, this study proposes forecasting future movement direction of a stock market index using its past returns along with those of another stock index and a currency. The second experiment is more favorable in terms of separability, as it forecasts the present movement direction of a stock market index using the present returns of indexes having many components in common. By construction, the latter experiment generally leads to higher performance metrics as the classes are more clearly split. However, using less easily separable and noisy market finance data, this study contributes to expending our experimental knowledge.

This paper ultimately proposes a forward-looking research direction to evaluate whether financial practitioners can already rely on quantum-enhanced SVM, for preparing a potential transition. Based on the findings, the classical version continues to exhibit higher accuracy compared to the quantum-enhanced counterpart. It is worth noting that the quantum metrics still fall within an acceptable range. While the results maintain a consistent average performance, there is some variability across different attempts that warrants attention. Furthermore, the duration of execution can vary significantly. Given these nuances, practitioners should remain cautious about immediate adoption while proactively empowering themselves to harness these technologies effectively.

Section 2 introduces quantum computing and the applied quantum machine learning literature. Section 3 presents the data set and briefly explains the SVM and the quantum-enhanced SVM algorithms. Then, the results, on both future and present forecasts of the movement direction of stock market indexes, are discussed in section 4. Finally, section 5 concludes this study and suggests some prospective research potentials.

2 Literature review

Quantum computers remain sensitive to noise so that small changes in temperature or vibration can make the qubits decohere. The technique to compensate for these imperfections is called quantum error correction. For decades, there has been ongoing academic research on quantum error correction (Cory et al., 1998; Knill et al., 2000; Chiaverini et al., 2004) and it remains highly active at present (Krinner et al., 2022; Smith et al., 2023; Sivak et al., 2023). Researchers suspect that reaching quantum advantage with uncorrected errors is impossible. But they remain unable to prove it for all cases. Aharonov et al. (2023) have taken a major step toward comprehensive proof that error correction is necessary for a lasting quantum advantage in random circuit sampling. And quantum computing companies are making slow but steady progress toward reducing the noise on both the hardware and software levels. For instance, IBM's 65-qubit systems from 2020 show twice the coherence time compared to when they first launched. IBM also considers an approach that is more in line with near-term goals using available hardware, the error mitigation techniques. The system is regularly checked for noise and then those noisy circuits are inverted to enable the generation of virtually error-free results. According to Schuld and Petruccione (2018) and Park et al. (2020b), we currently are in the "noisy intermediate-scale quantum" (NISQ) era. NISQ processors are prone to decoherence and are not able to continuously correct quantum errors.

Regarding quantum machine learning, Biamonte et al. (2017) define it as the exploration of "how to devise and implement quantum software that could enable machine learning that is faster than that of classical computers." But they concede that both the hardware and software challenges remain considerable. With more specificities, Liu and Rebentrost (2018) propose two widely used machine learning algorithms' quantum versions (i.e., kernel principal component analysis and one-class SVM) to detect outliers in quantum states. For instance, a set-up that can be applied in finance, fraud detection[8], and surveillance. Martin et al. (2021) use quantum principal component analysis to decrease the number of noisy factors needed to simulate the time evolution of diverse time-maturing forward rates. Additionally, a quantum

[8]See Cecchini et al. (2010) for a classical version.

version of the widely used k-means clustering was developed (Lloyd et al., 2013; Kerenidis et al., 2019). Emmanoulopoulos and Dimoska (2022) explore the efficacy of quantum neural networks for forecasting time series signals with simulated quantum forward propagation.

The SVM algorithm reaches limitations when the feature space becomes large and the kernel functions become computationally costly to estimate. That is why quantum algorithms' computational speed-ups is important. Rebentrost et al. (2014) show that the SVM can be implemented on a quantum computer and can achieve an exponential speed-up. One strong theoretical point is that the time complexity of the SVM is $\mathcal{O}(N)$ and is scaled down to $\mathcal{O}(\log N)$ with the quantum version (Lloyd et al., 2013; Rebentrost et al., 2014; Li et al., 2015). They show that quantum SVM can solve certain classification problems which classical computers cannot efficiently handle. Other recent theoretical papers implement the SVM on quantum computers (Chatterjee and Yu, 2017; Schuld and Killoran, 2019; Havlicek et al., 2019; Park et al., 2020a; Liu et al., 2021). However, exponential speed-up is attainable through the computations, thus only if data is provided in a coherent superposition. It is no longer possible if they are brought conventionally from a classical computer. My experiments are based on actual financial data that is not generated by a quantum computer; an issue overcome by Havlicek et al. (2019). Inputs are converted from classical bits to quantum bits (qubits) before the quantum algorithm is run. Once the computations are finished, outputs' qubits are turned back into bits. In line with McRae and Hilke (2020), it is important to be aware that the quantum SVM version running on Qiskit is not entirely quantum but rather "quantum-enhanced" as the quantum processor performs only certain operations, the rest are done classically. Havlicek et al. (2019) define two methods that map non-linearly classically provided data to a quantum state. The first method, a major alternative to the quantum SVM, called variational quantum classifier (VQC), uses a variational quantum circuit to classify the data as a conventional SVM does. The second method is the quantum kernel estimator (QKE) that uses the quantum processor to estimate the kernel function and then optimizes a classical SVM.

The established research has presented various ways of relating finance and quantum computing. For instance, Rebentrost and Lloyd (2018), focusing on quantum speed-ups, proposed an algorithm that comes up with the optimal risk-return tradeoff curve based on quantum access to past returns. Rebentrost et al. (2018) develop a quantum algorithm for the Monte Carlo pricing of financial derivatives. Notably, they handle the preparation of the relevant probability distributions in a quantum superposition, the implementation of the payoff functions through quantum circuits, and the extraction of financial derivatives' prices through quantum measurements. Stamatopoulos et al. (2020) develop a methodology to price options and portfolios of options using amplitude estimation. Their algorithm provides a quadratic speed-up compared to classical Monte Carlo methods. A global discussion on the application of quantum computing to financial issues is brought up by Orùs et al. (2019). They deal with various subjects like portfolios optimization, seeking arbitrage opportunities, credit scoring, and how quantum amplitude estimation can lead to quantum speed-ups for Monte Carlo sampling that has applications such as derivatives pricing and risk analysis. Another example is the mapping of the Black-Scholes-Merton formula (Black and Scholes, 1973; Merton, 1973) to Schrödinger equation (Haven, 2002). This emphasizes that finance could benefit from a computational speed-up which "could manifest itself in a number of different ways, each of which could imply gargantuan savings for governments, financial institutions, and individuals" (Orùs et al., 2019). Egger et al. (2020) focus on applications across asset management, investment banking, retail and corporate banking involving simulation, optimization, and machine-learning problems, where quantum computing provides more reliable solutions.

Large-scale quantum computing remains theoretical, but cloud-based technologies allow for small-scale empirical experiments. The quantum-enhanced SVM has mainly been tested on separable data sets, such as "Breast Cancer Wisconsin (Diagnostic)," "Iris" or "Wine" data sets (Havenstein et al., 2018; Sarma et al., 2019; McRae and Hilke, 2020), or to solve a minimal optical character recognition (OCR) problem in Li et al. (2015). In addition to dealing with "Breast Cancer" and "Wine" UCI data sets, McRae and Hilke (2020) generate three classes of data: ad-hoc data, Anderson data (Anderson, 1958), and Coronavirus data (Smith et al., 2020). Underlining the current involvement in applied quantum SVM research, Peters et al. (2021) apply the quantum kernel SVM to a cosmological benchmark using real spectral features. But to the best of my knowledge, noisier and non-stationary data sets, such as market finance time series, were not handled in previous studies.

The first analysis was inspired by the framework of Huang et al. (2005), which used the SVM to forecast the movement direction of a stock market index. Forecasting stock returns' direction using machine learning algorithms and, more precisely, the SVM, remains popular (Gong et al., 2016; Chen and Hao, 2017; Henrique et al., 2018; Ismail et al., 2020; Akyildirim et al., 2022). As an attempt to provide more empirical experiences to bridge the research gap between the application of the SVM and the quantum-enhanced SVM, this study assesses how the quantum-enhanced SVM behaves compared to the SVM in forecasting stock market indexes' up and down movements.

3 Methodology

3.1 Model formulation

The usage of cloud-based quantum devices constrains the experiments to limit both the dimension and the size of the data set. Huang et al. (2005) forecast the movement direction of a stock market index with only two features. Running a conventional SVM, the authors predict the Japanese NIKKEI 225 index using the S&P 500 index (hereafter abbreviated to SPX index), standing as a well-known indicator of the American economic condition, and the yen, affecting the Japanese export. In this study, I set the Canadian S&P/TSX composite index (hereafter abbreviated to TSX index) as a dependent variable. Notably, China became the most significant export target for Japan[9] and the yuan is restricted to fluctuating between bands (Jermann et al., 2022). This study favors the TSX index as Canada now has many similarities with Japan then: being a major exporter, predominantly exporting to the United States[10], and belonging to the high-GDP countries. Additionally, Huang et al. (2005) noticed that the "Japanese interest rate has dropped down to almost zero since 1990," a comparable situation to the Canadian interest rate during the period studied (cf. fig. 4 in the Appendix). The TSX index is here explained with three features: its own past values and those of the SPX index and the Canadian/US dollar exchange rate (hereafter CAD).

According to Timmermann and Granger (2004) the efficient-market hypothesis does not rule out forecasting short-lived gains. But, with such a small sample it may not be possible to efficiently compare the models. Therefore, in a second experiment, I run the algorithms on a more favorable framework, predicting present movements of the Euro Stoxx 50 index (hereafter abbreviated to SX5E index) with highly correlated indexes. Namely, the German and French indexes (DAX and CAC indexes), respectively correlated to the SX5E index, with which they share the most important components, by 98.2% and 95%. By construction, this experiment should lead to higher performance metrics as the classes are more clearly split.

For both settings, the data set spans weekly for more than ten years, between two major crises (i.e., the GFC and the COVID pandemic), from 2009-07-08 until 2020-01-22, and is

[9] https://oec.world/en/profile/country/jpn

[10] Around 75% of Canada's exports go to the United States (https://oec.world/en/profile/country/can).

based entirely on Bloomberg L. P. data. To avoid being disturbed by the potential end-of-week volatility, my experiment was done on Wednesdays, as in Huang et al. (2005). The training sets show a slight imbalance[11], while the test sets sometimes exhibit a clear imbalance. Consequently, reliance on the accuracy metric is not assured. To address this, I evaluate the balanced accuracy, F1-score, and the components of the latter: precision and recall.

In this study, the "SVC" class of the machine learning library called "sci-kit learn" (Pedregosa et al., 2011) is used to implement the classical support vector classification on a local computer. And a dedicated NISQ version of the quantum-enhanced SVM algorithm is run on a freely publicly available cloud-based quantum device. The free access is limited to 5 qubits and is provided by IBM Q Experience (IBM Staff, 2016). The program is written with the Python-based Quantum information science kit (Qiskit, IBM Staff (2017)). Aware of the cloud-based quantum technology's limitations I restrict, for both experiments, the size of the training and test sets to respectively 80 and 20 consecutive weeks. Both experiments are run monthly over the same 113 Wednesdays, providing 113 performance measures which are summarized and analyzed in section 4.

To put it in a nutshell, this study forecasts the one-week direction of the TSX composite index and the current value of the Euro Stoxx 50 index's direction. It can be expressed as follows:

$$Direction_{t+1}^{TSX} = f\left(Return_t^{TSX}, Return_t^{SPX}, Return_t^{CAD}\right) \tag{1}$$

$$Direction_t^{SX5E} = f\left(Return_t^{DAX}, Return_t^{CAC}\right) \tag{2}$$

3.2 Classical Support Vector Machine

This article deals with market finance time series that are implicitly noisy and non-stationary. According to Cao and Tay (2003), the SVM has successfully modelled financial time series. These authors also point out interesting characteristics of this algorithm such as good generalization performance, the absence of local minima, and sparse representation of solution.

In this paper, I determine the TSX and SX5E indexes' movement direction with binary classification. The algorithm will learn a discriminant rule (here, the maximum-margin hyperplane) from the training set to summarize the main idea quickly. In this framework, each example,

[11] Less than 60%, mainly in favor of up movements.

out of T (here 113 observations) training examples, is a vector that belongs to the vector space of D considered attributes (here stock indexes and forex variables), $\mathbf{x}_t \in \mathbb{R}^D$. In the training set, examples are labelled as the algorithm must know to which class (the explained indexes are either going up or down), $y_t \in \{-1, 1\}$, corresponds to each data point. Once the discriminant rule is known, the algorithm can handle a new example and determine to which class it most likely belongs.

Boser et al. (1992)' article and the soft margin improvement proposed by Cortes and Vapnik (1995) are the origin of the SVM that is essentially based on the blending of two key ideas: maximum-margin hyperplane and kernel function. The aim is to select the hyperplane for which its distance to the two closest points is maximized. This hyperplane is known as the maximum-margin hyperplane. The primal problem with soft margin, writing the regularization term C, the slack variable ϵ_t and the weights plus the bias $(\mathbf{w}, b) \in \mathbb{R}^{D+1}$, is commonly stated as follows:

$$
\begin{aligned}
\min_{\mathbf{w}} \quad & \frac{1}{2}\|\mathbf{w}\|^2 + C\sum_{t=1}^{T} \epsilon_t \\
\text{subject to} \quad & y_t(\mathbf{w}^\top \mathbf{x}_t + b) \geq 1 - \epsilon_t, \\
& C \geq 0, \epsilon_t \geq 0.
\end{aligned}
$$

Although such a problem has a straightforward solution for linear classification, it is not anymore the case for non-linear classification. But one can consider mapping the data into a higher feature space of dimension D', such that $D' > D$. Let's define $\Phi(\mathbf{x})$ as the feature map embedding the data to the larger space, a linear classifier can then be rewritten as $h(\mathbf{x}) = \mathbf{w}^\top \Phi(\mathbf{x}) + b$. However, calculations to map all data in a higher feature space become computationally costly as the number of dimensions increases. But it is enough for one to use a kernel function to avoid long computation in high dimensions. The idea is to act as if the data is mapped to a higher-dimensional space, in which a linear classifier could be used as a maximum-margin hyperplane. This method, called the "kernel trick," allows operating in the original feature space by applying the kernel function to the data. In other words, instead of computing the inner products of the images of each datum in higher dimension $\Phi(\mathbf{x}_i)^\top \Phi(\mathbf{x}_j)$, one can use the kernel function to directly return the inner product of the transformed data in the higher feature space without any actual computation of $\Phi(\mathbf{x})$, such that $K(\mathbf{x}_i, \mathbf{x}_j) = \Phi(\mathbf{x}_i)^\top \Phi(\mathbf{x}_j)$. To make a connection with the following quantum version, I write the dual problem maximizing the Lagrangian:

$$\begin{aligned} &\text{Maximize } \tilde{L}(\alpha) = \sum_{t=1}^{T} \alpha_t - \frac{1}{2}\sum_{i=1}^{T}\sum_{j=1}^{T} \alpha_i \alpha_j y_i y_j K(\mathbf{x}_i, \mathbf{x}_j) \\ &\text{subject to } \sum_{t=1}^{T} \alpha_t y_t = 0, \text{ and } \alpha_t \geq 0 \text{ for each } t. \end{aligned} \tag{3}$$

Thus, the hyperplane can be rewritten $h(\mathbf{x}) = \sum_{t=1}^{T} \alpha_t^* y_t K(\mathbf{x}_t, \mathbf{x}) + b$, where α^* is the optimal Lagrange multiplier. In this study, I run the classical SVM and I select the radial basis function ($K(\mathbf{x}_i, \mathbf{x}_j) = \exp\{-\gamma \mid\mid \mathbf{x}_i - \mathbf{x}_j \mid\mid^2\}$) as the kernel function. I also use the linear kernel ($K(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{x}_i^\top \mathbf{x}_j$) as it does not require to optimize the γ parameter.

3.3 Quantum-enhanced Support Vector Machine

The experiments also use the quantum kernel estimator (QKE), a method mapping non-linearly classically provided data to a quantum state (Havlicek et al., 2019). It estimates the kernel function with the quantum processor and then optimizes an SVM. In other words, $K(\mathbf{x}_i, \mathbf{x}_j)$ is estimated quantumly and then the dual problem (cf. eq. (3)) is classically solved. This algorithm sometimes called SVM-QKE (Liu et al., 2021) is available in Qiskit (Qiskit Development Team, 2020).

As already stated above, the non-linear transformation function applied to the data that cannot be linearly separated in their original space is called a feature map. In this new feature space, one classifies by assessing the examples' closeness, corresponding to computing each pair of data inner product. But calculating the non-linear feature map for each datum is not necessary. It is sufficient to evaluate each pair of examples' inner product in the new feature space. This set of inner products is simpler to compute and is named the kernel. Quantum-enhanced SVM algorithms are used when the required feature map is not efficient classically, notably if the necessary computational resources are expected to scale exponentially with the problem size. The QKE uses the quantum processor to estimate the kernel in the feature space to solve this dilemma. In this set-up, the quantum computer is used twice, to calculate the kernel for all pairs of data from the training set and then estimate the kernel for a new datum from the test set.

More explicitly, the kernel's inputs are the fidelities, measures of quantum states' closeness, between the feature vectors. The standard method to estimate fidelities is via the swap test

(Buhrman et al., 2001), but as only the fidelities' values are needed, Havlicek et al. (2019) rather choose a shallower circuit as proposed by Cincio et al. (2018). In this case, the states in the feature map are structured allowing to estimate of the overlap from the transition amplitude, the kernel is thus set as $K(\mathbf{x}_i, \mathbf{x}_j) = |\langle\Phi(\mathbf{x}_i)|\Phi(\mathbf{x}_j)\rangle|^2 = |\langle 0| U^{\dagger}_{\Phi(\mathbf{x}_i)} U_{\Phi(\mathbf{x}_j)} |0\rangle|^2$. In other words, the quantum circuit $U^{\dagger}_{\Phi(\mathbf{x}_i)} U_{\Phi(\mathbf{x}_j)}$ is applied to the initial reference state $|0\rangle$. Then, the resulting state, $U^{\dagger}_{\Phi(\mathbf{x}_i)} U_{\Phi(\mathbf{x}_j)} |0\rangle$, is sampled R (number of shots[12]) times in the z-basis. Finally, the kernel estimator, up to a sampling error $\widetilde{\epsilon} = \mathcal{O}(R^{-0.5})$, is $\#\{0, \dots, 0\}/R$ that corresponds to the number of observed zero bit-strings divided by R. Once the kernel matrix has been built for the entire training set, the separating hyperplane is found by classically solving eq. (3). Then, the quantum computer is called again to estimate the kernel for a new datum $\mathbf{x}' \in \mathbb{R}^D$ with all the support vectors, and the new corresponding label is assigned by $\hat{y}' = \text{sign}\big(\sum_{t=1}^{T} \alpha_t y_t K(\mathbf{x}_t, \mathbf{x}') + b\big)$.

From a hardware viewpoint, in the manner of elementary logic gates of a classical computer, quantum computers handle a series of quantum gates[13] (called quantum circuits), which were particularity inherited from quantum properties to be reversible. For eq. (1) and eq. (2), the chosen features map is a second-order Pauli-Z evolution circuit (cf. a simplified example in fig. 1) with a number of qubits, $n_{qb} = 3$ and $n_{qb} = 2$[14], respectively — the quantum computer manipulates $2^{n_{qb}}$ probability values — two repetitions and full entanglement.

[Insert Figure 1]

From a practical point of view, it is interesting to emphasize that SVM algorithms require choosing a kernel out of several kernel functions, e.g., polynomial, Gaussian, radial basis, sigmoid, etc. It also requires tuning its hyperparameters accordingly to optimize the bias-variance tradeoff. Notably, non-linear kernel functions require to set the γ parameter. This is not the case for the quantum SVM by Havlicek et al. (2019) that is only derived from $K(\mathbf{x}_i, \mathbf{x}_j) = |\langle\Phi(\mathbf{x}_i)|\Phi(\mathbf{x}_j)\rangle|^2$, making it simpler for anyone to use. However, for a specific training set, with the SVM, one will always obtain the same output[15], which is untrue for the inherently stochastic quantum-enhanced SVM.

In the following section 4, I present the empirical results of the comparisons between the SVM and the quantum-enhanced SVM on actual financial data, both in terms of accuracy and

[12] In this experiment $R = 2^{10}$.
[13] The quantum gates used in this study are available in the Appendix, and more details are available in Kurowski et al. (2023).
[14] n_{qb} corresponds to the number of features.
[15] Assuming that no random state parameter is tuned.

computation time.

4 Results

4.1 Forecasting the future movement direction of stock market index

Equity forecasting remains of utmost importance for financial investors (Timmermann and Granger, 2004; Pástor and Stambaugh, 2009; Rapach et al., 2013; Nyberg and Pönkä, 2016). In the first experiment, the TSX index movement direction is predicted using its own returns of the previous weeks, with those of the SPX index and the CAD. As shown in fig. 2, there is no clear separation of the distribution of the movement directions. Thus, predicting where the market is heading is challenging by construction.

[Insert Figure 2]

[Insert Table 1]

Numerical outputs from table 1 match the lack of separability displayed on fig. 2. Although the radial basis function (RBF) kernel versions provide the best balanced accuracies, they remain only slightly above 50%. The linear and quantum kernels, that do not require a γ parameter, offer similar results, but closer to the 50% threshold. In other words, forecasting future weekly movement direction of the TSX index using a selected limited number of features with classical and quantum SVM does not seem appropriate. Obviously, with around one chance out of two to correctly predict, no investment decision should be based on these frameworks.

Although the F1-Score is relatively high, I mainly note that recall is higher than precision, meaning that out of the up movements that should have been selected many were actually selected. Yet the lower precision highlights that only some of the up movements retrieved are actually relevant.

Under current conditions, it is not beneficial to employ these algorithms on a limited set of features to predict future movement direction of the TSX index. However, these results may be overcome with the possibility of performing the same computations on bigger matrices, notably with the improvement of quantum technologies' capacities.

In section 4.2, I perform the same analysis on a more favorable set-up. The chosen framework has the advantage of splitting the classes more clearly, allowing to study of the algorithms' behavior in better conditions. More precisely, I propose to forecast the present movement direction of an index (cf. section 3.1). From a practitioner's point of view, it can be seen as a missing data completion task.

4.2 Forecasting the present movement direction of stock market index

In this section, up and down weekly movement directions of the SX5E index are forecasted with the returns of the DAX and CAC indexes occurring on the same dates. This set-up ensures algorithm comparability, albeit in a basic framework. However, it is worth noting that predictability might be absent in section 4.1, potentially leading to non-comparability of the models. By construction, fig. 3 shows a clearer separation of the distribution of these movement directions than in fig. 2.

[Insert Figure 3]

[Insert Table 2]

The good separability displayed in fig. 3 is reflected in the numerical outputs. The best balanced accuracies are found for each algorithm using the highest tested value of C. In other terms, when the strength of the regularization is the weakest, it leads to the highest number of well-classified training set points but potentially inclines to overfit. However, the results presented in table 2 are computed on the test set and thus, there is no overfitting noted. Regarding the SVM with RBF kernel, large values of γ allowing "more" non-linearity in the hyper-plane also leads to higher accuracies. As in section 4.1 the SVM with RBF kernel provides the best balanced accuracies, above 90%. The easier-to-tune linear and quantum kernels lead to lower balanced accuracies but still higher than 80%. Overall, for such a task, a financial practitioner can still improve the accuracy, in this experiment by more than 8%, preferring the classical non-linear kernel SVM to the quantum-enhanced SVM.

In this experiment, the F1-Score is relatively high and both recall and precision provide high figures. Thus, among the up movements that should have been selected, many were actually selected, and most of the up movements retrieved are actually relevant.

Due to the better metrics obtained from the SVM with the RBF kernel, the current advice is against utilizing the quantum-enhanced version for forecasting the present movement direction of the SX5E index. It is worth noting, however, that the quantum-enhanced version maintains the advantage of being easier to tune, which in turn reduces the likelihood of parametrization mistakes. Again, these results may be overcome with the possibility of performing the same computations on bigger matrices, notably with the improvement of quantum technologies' capacities.

4.3 Technical considerations

In this analysis, it is important to note that while the results may vary between individual runs, at the example scale, they demonstrate a level of stability on average, instilling confidence in their overall reliability. Besides, the quantum-enhanced SVM's average computation time per observation using the cloud-based quantum computer is 3 hours and 55 minutes, — ranging from 40 minutes to more than 8 hours, making the computation duration unreliable —, being on average 5.9×10^6 times slower than performing classical computation locally. McRae and Hilke (2020) also using Qiskit's quantum-enhanced SVM, reached similar conclusions, both in terms of accuracy and calculation time, and discussed further the difference between the time of computation in practice from the theoretical exponential speed-up.

In sum, the quantum-enhanced SVM does not demonstrate better performance, in terms of both accuracy and computation time, compared to the classical SVM.

5 Conclusion

In this paper, two empirical experiments compare the predictive powers of the quantum-enhanced SVM and the classical SVM based on stock market index data, which provide preliminary outcomes and financial insights for academics and practitioners. The quantum algorithm is run on an actual quantum computer accessible through a cloud-based technology. However, the main limitation of this technology lies in the fact that these freely publicly available quantum devices can only deal with small input matrices within reasonable computation time.

The main result is that the quantum-enhanced SVM underperforms the SVM in terms of both accuracy and computation time. While the current indication suggests that financial professionals may not immediately gain from the quantum-enhanced SVM for forecasting stock market index movement direction, it is possible that in the future, there could be an opportunity. Over time, researchers and practitioners might discover the advantage in favoring the quantum SVM over the classical SVM. Beyond its potential to provide significant speed-up, the quantum SVM eliminates the need for selecting a kernel function and fine-tuning a large set of hyperparameters.

Embracing the trend of developing quantum technologies, which encompasses quantum sensors, simulators, communications, and computers, this paper directs its attention to a financial study, centering on quantum computers. The contribution of this paper is two-fold. First, this paper synthesizes the quantum machine learning literature with a focus on quantum SVM to academics and practitioners in the realm of finance. Second, this paper conducts two pilot experiments using the quantum-enhanced SVM, grasping the current state of progress of freely publicly available quantum technologies, which proposes a future research agenda. Quantum computing continuously faces various challenges in its development, due to the limited number of qubits available, the short amount of time to quantum decoherence, the required frigid temperatures to operate, etc. Hence, the outputs could be improved in future research, once a more efficient technology is available.

6 Appendix

6.1 Bank of Canada overnight lending rate

[Insert Figure 4]

6.2 Quantum gates

The second-order Pauli-Z evolution circuit relies the following gates:

- The Hadamard gate which acts on a single qubit and creates superposition.
- The controlled-NOT gate which entangles and disentangles qubits together (Monroe et al., 1995).
- The U1 gate that is a diagonal gate leading to a single rotation about the z-axis (McKay et al., 2017).

7 Acknowledgments

I am most grateful to E. Combarro, M. Tod, and A. Bignú for their detailed technical advice, and to L. Xu for editing and to J. P. Lardy proofreading this article.

References

Aharonov, D., Gao, X., Landau, Z., Liu, Y., and Vazirani, U. (2023). A polynomial-time classical algorithm for noisy random circuit sampling. *arXiv:2211.03999*, pages 1–27.

Akyildirim, E., Bariviera, A. F., Nguyen, D. K., and Sensoy, A. (2022). Forecasting high-frequency stock returns: a comparison of alternative methods. *Annals of Operations Research*, pages 1–52.

Alcazar, J., Leyton-Ortega, V., and Perdomo-Ortiz, A. (2020). Classical versus quantum models in machine learning: insights from a finance application. *Machine Learning: Science and Technology*, 1(035003):1–10.

Anderson, P. W. (1958). Absence of diffusion in certain random lattices. *Physical Review*, 109(5):1492–1505.

Arute, F., Arya, K., Babbush, R., Bacon, D., Bardin, J. C., Barends, R., Biswas, R., Boixo, S., Brandao, F. G. S. L., Buell, D. A., Burkett, B., Chen, Y., Chen, Z., Chiaro, B., Collins, R., Courtney, W., Dunsworth, A., Farhi, E., Foxen, B., Fowler, A., Gidney, C., Giustina, M., Graff, R., Guerin, K., Habegger, S., Harrigan, M. P., Hartmann, M. J., Ho, A., Hoffmann, M., Huang, T., Humble, T. S., Isakov, S. V., Jeffrey, E., Jiang, Z., Kafri, D., Kechedzhi, K., Kelly, J., Klimov, P. V., Knysh, S., Korotkov, A., Kostritsa, F., Landhuis, D., Lindmark, M., Lucero, E., Lyakh, D., Mandra, S., McClean, J. R., McEwen, M., Megrant, A., Mi, X., Michielsen, K., Mohseni, M., Mutus, J., Naaman, O., Neeley, M., Neill, C., Niu, M. Y., Ostby, E., Petukhov, A., Platt, J. C., Quintana, C., Rieffel, E. G., Roushan, P., Rubin, N. C., Sank, D., Satzinger, K. J., Smelyanskiy, V., Sung, K. J., Trevithick, M. D., Vainsencher, A., Villalonga, B., White, T., Yao, Z. J., Yeh, P., Zalcman, A., Neven, H., and Martinis, J. M. (2019). Quantum supremacy using a programmable superconducting processor. *Nature*, 574:505–510.

Bas, J., Cirillo, C., and Cherchi, E. (2021). Classification of potential electric vehicle purchasers: A machine learning approach. *Technological Forecasting and Social Change*, 168:120759–(1–13).

Begušió, T. and Chan, G. K.-L. (2023). Fast classical simulation of evidence for the utility of quantum computing before fault tolerance. *arXiv:2306.16372*, pages 1–4.

Ben-Tal, A., Hazan, E., Koren, T., and Mannor, S. (2015). Oracle-based robust optimization via online learning. *Operations Research*, 63(3):628–638.

Benítez-Peña, S., Blanquero, R., Carrizosa, E., and Ramírez-Cobo, P. (2023). Cost-sensitive probabilistic predictions for support vector machines. *European Journal of Operational Research*, in Press:1–12.

Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., and Lloyd, S. (2017). Quantum machine learning. *Nature*, 549:195–202.

Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654.

Boser, B., Guyon, I. M., and Vapnik, V. (1992). A training algorithm for optimal margin classifier. *Proceedings of the Fifth Annual ACM Workshop on Computational Learning Theory*, 5:144–152.

Buhrman, H., Cleve, R., Watrous, J., and de Wolf, R. (2001). Quantum fingerprinting. *Physical Review Letters*, 87(16):167902–(1–4).

Cao, L. J. and Tay, F. E. H. (2003). Support vector machine with adaptive parameters in financial time series forecasting. *IEEE Transactions on Neural Networks*, 14(6):1506 – 1518.

Cecchini, M., Aytug, H., Koehler, G. J., and Pathak, P. (2010). Detecting management fraud in public companies. *Management Science*, 56(7):1146–1160.

Chatterjee, R. and Yu, T. (2017). Generalized coherent states, reproducing kernels, and quantum support vector machines. *Quantum Information & Computation*, 17:1292–1306.

Chen, Y. and Hao, Y. (2017). A feature weighted support vector machine and k-nearest neighbor algorithm for stock market indices prediction. *Expert Systems with Applications*, 80:340–355.

Chiaverini, J., Leibfried, D., Schaetz, T., Barrett, M. D., Blakestad, R. B., Britton, J., Itano, W. M., Jost, J. D., Knill, E., Langer, C., Ozeri, R., and Wineland, D. J. (2004). Realization of quantum error correction. *Nature*, 432:602–605.

Cincio, L., Subasi, Y., Sornborger, A. T., and Coles, P. J. (2018). A rigorous and robust quantum speed-up in supervised machine learning. *arXiv:1803.04114*, pages 1–12.

Cortes, C. and Vapnik, V. (1995). Support-vector networks. *Machine Learning*, 20(3):273–297.

Cory, D. G., Price, M. D., Maas, W., Knill, E., Laflamme, R., Zurek, W. H., Havel, T. F., and Somaroo, S. S. (1998). Experimental quantum error correction. *Physical Review Letters*, 81(10):2152–2155.

Egger, D. J., Gambella, C., Marecek, J., McFaddin, S., Mevissen, M., Raymond, R., Simonetto, A., Woerner, S., and Yndurain, E. (2020). Quantum computing for finance: State-of-the-art and future prospects. *IEEE Transactions on Quantum Engineering*, 1:3101724(1–24).

Emmanoulopoulos, D. and Dimoska, S. (2022). Quantum machine learning in finance: Time series forecasting. *arXiv:2202.00599*, pages 1–20.

Gong, X., Si, Y.-W., Fong, S., and Biuk-Aghai, R. P. (2016). Financial time series pattern matching with extended ucr suite and support vector machine. *Expert Systems with Applications*, 55:284–296.

Harper, R. and Flammia, S. T. (2019). Fault-tolerant logical gates in the ibm quantum experience. *Physical Review Letters*, 122(8):080504–(1–6).

Haven, E. E. (2002). A discussion on embedding the black–scholes option pricing model in a quantum physics setting. *Physica A: Statistical Mechanics and its Applications*, 304(3-4):507–524.

Havenstein, C., Thomas, D., and Chandrasekaran, S. (2018). Comparisons of performance between quantum and classical machine learning. *SMU Data Science Review*, 1(4-11):1–29.

Havlicek, V., Corcoles, A. D., Temme, K., Harrow, A. W., Kandala, A., Chow, J. M., and Gambetta, J. M. (2019). Supervised learning with quantum-enhanced feature spaces. *Nature*, 567:209–212.

Hebenstreit, M., Alsina, D., Latorre, J. I., and Kraus, B. (2017). Compressed quantum computation using a remote five-qubit quantum computer. *Physical Review A*, 95(5):052339–(1–7).

Henrique, B. M., Sobreiro, V. A., and Kimura, H. (2018). Stock price prediction using support vector regression on daily and up to the minute prices. *The Journal of Finance and Data Science*, 4(3):183–201.

Huang, W., Nakamoria, Y., and Wang, S.-Y. (2005). Forecasting stock market movement direction with support vector machine. *Computers & Operations Research*, 32(10):2513–2522.

IBM Staff (2016). IBM quantum. https://quantum-computing.ibm.com/. [Online; accessed 4-May-2016].

IBM Staff (2017). Qiskit: Quantum information science kit. https://qiskit.org/. [Online; accessed 7-March-2017].

Ismail, M. S., Noorani, M. S. M., Ismail, M., Razak, F. A., and Alias, M. A. (2020). Predicting next day direction of stock price movement using machine learning methods with persistent homology: Evidence from kuala lumpur stock exchange. *Applied Soft Computing*, 93(106422):1633–1662.

Jermann, U. J., Wei, B., and Yue, V. Z. (2022). The two-pillar policy for the rmb. *The Journal of Finance*, 77(6):3093–3140.

Jiménez-Cordero, A., Morales, J. M., and Pineda, S. (2021). A novel embedded min-max approach for feature selection in nonlinear support vector machine classification. *European Journal of Operational Research*, 293(1):24–35.

Kang, H., Zong, X., Wang, J., and Chen, H. (2023). Binary gravity search algorithm and support vector machine for forecasting and trading stock indices. *International Review of Economics & Finance*, 84:507–526.

Kerenidis, I., Landman, J., Luongo, A., and Prakashs, A. (2019). q-means: A quantum algorithm for unsupervised machine learning. *Advances in Neural Information Processing Systems 32*.

Kim, Y., Eddins, A., Anand, S., Wei, K. X., Van den Berg, E., Rosenblatt, S., Nayfeh, H., Wu, Y., Zaletel, M., Temme, K., and Kandala, A. (2023). Evidence for the utility of quantum computing before fault tolerance. *Nature*, 618:500–505.

Knill, E., Laflamme, R., and Viola, L. (2000). Theory of quantum error correction for general noise. *Physical Review Letters*, 84(11):2525–2528.

Krinner, S., Lacroix, N., Remm, A., Paolo, A. D., Genois, E., Leroux, C., Hellings, C., Lazar, S., Swiadek, F., Herrmann, J., Norris, G. J., Andersen, C. K., Müller, M., Blais, A., Eichler, C., and Wallraff, A. (2022). Realizing repeated quantum error correction in a distance-three surface code. *Nature*, 605:669–674.

Kurowski, K., Pecyna, T., Slysz, M., Rózycki, R., Waligóra, G., and Węglarz, J. (2023). Application of quantum approximate optimization algorithm to job shop scheduling problem. *European Journal of Operational Research*, 310(2):518–528.

Li, Z., Liu, X., Xu, N., and Du, J. (2015). Experimental realization of a quantum support vector machine. *Physical Review Letters*, 114(14):140504–(1–5).

Liu, N. and Rebentrost, P. (2018). Quantum machine learning for quantum anomaly detection. *Physical Review A*, 97(4):042315–(1–10).

Liu, Y., Arunachalam, S., and Temme, K. (2021). A rigorous and robust quantum speed-up in supervised machine learning. *Nature Physics*.

Lloyd, S., Mohseni, M., and Rebentrost, P. (2013). Quantum algorithms for supervised and unsupervised machine learning. *arXiv:1307.0411*, pages 1–11.

Luo, J., Hong, T., Gao, Z., and Fang, S.-C. (2023). A robust support vector regression model for electric load forecasting. *International Journal of Forecasting*, 39(2):1005–1020.

Madsen, L. S., Laudenbach, F., Falamarzi-Askarani, M., Rortais, F., Vincent, T., Bulmer, J. F. F., Miatto, F. M., Neuhaus, L., Helt, L. G., Collins, M. J., Lita, A. E., Gerrits, T., Nam, S. W., Vaidya, V. D., Menotti, M., Dhand, I., Vernon, Z., Quesada, N., and Lavoie, J. (2022). Quantum computational advantage with a programmable photonic processor. *Nature*, 606:75–81.

Maldonado, S., Pérez, J., and Bravo, C. (2017). Cost-based feature selection for support vector machines: An application in credit scoring. *European Journal of Operational Research*, 261(2):656–665.

Martin, A., Candelas, B., Rodríguez-Rozas, A., Martín-Guerrero, J. D., Chen, X., Lamata, L., Orús, R., Solano, E., and Sanz, M. (2021). Toward pricing financial derivatives with an ibm quantum computer. *Physical Review Research*, 3(1):013167–(1–12).

McKay, D. C., Wood, C. J., Sheldon, S., Chow, J. M., and Gambetta, J. M. (2017). Efficient Z gates for quantum computing. *Physical Review A*, 96(2):022330–(1–8).

McRae, P.-A. and Hilke, M. (2020). Quantum-enhanced machine learning for covid-19 and anderson insulator predictions. *arXiv:2012.03472*, pages 1–25.

Merton, R. C. (1973). Theory of rational option pricing. *The Bell Journal of Economics and Management Science*, 4(1):141–183.

Monroe, C., Meekhof, D. M., King, B., Itano, W. M., and Wineland, D. J. (1995). Demonstration of a fundamental quantum logic gate. *Physical Review Letters*, 75(25):4714–4717.

Montanaro, A. (2015). Quantum speedup of Monte Carlo methods. *Proceedings of the Royal Society A*, 471(2181).

Mugel, S., Kuchkovsky, C., Sánchez, E., Fernández-Lorenzo, S., Luis-Hita, J., Lizaso, E., and Orús, R. (2022). Dynamic portfolio optimization with real datasets using quantum processors and quantum-inspired tensor networks. *Physical Review Research*, 4(1):013006–(1–12).

Nurdin, H. I. (2019). Quantum stochastic processes and the modelling of quantum noise. *Springer London, London*, page 1–8.

Nyberg, H. and Pönkä, H. (2016). International sign predictability of stock returns: The role of the united states. *Economic Modelling*, 58:323–338.

Orùs, R., Mugel, S., and Lizaso, E. (2019). Quantum computing for finance: Overview and prospects. *Reviews in Physics*, 4(100028):1–12.

Park, D. K., Blank, C., and Petruccione, F. (2020a). The theory of the quantum kernel-based binary classifier. *Physics Letters A*, 384(21):126422–(1–12).

Park, J.-E., Quanz, B., Wood, S., Higgins, H., and Harishankar, R. (2020b). Practical application improvement to quantum svm: theory to practice. *arXiv:2012.07725*, pages 1–9.

Pástor, L. and Stambaugh, R. F. (2009). Predictive systems: Living with imperfect predictors. *The Journal of Finance*, 64(4):1583–1628.

Pednault, E., Gunnels, J., Maslov, D., and Gambetta, J. (2019). On quantum supremacy. *IBM research blog*, page 1.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12:2825–2830.

Peters, E., Caldeira, J., Ho, A., Leichenauer, S., Mohseni, M., Neven, H., Spentzouris, P., Strain, D., and Perdue, G. N. (2021). Machine learning of high dimensional data on a noisy quantum processor. *npj Quantum Information*, 7(161):1–5.

Pistoia, M., Ahmad, S. F., Ajagekar, A., Buts, A., Chakrabarti, S., Herman, D., Hu, S., Jena, A., Minssen, P., Niroula, P., Rattew, A., Sun, Y., and Yalovetzky, R. (2021). Quantum machine learning for finance, ICCAD special session paper. *2021 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, pages 1–9.

Preskill, J. (2012). Quantum computing and the entanglement frontier. *arXiv:1203.5813*, pages 1–19.

Qiskit Development Team (2020). Qiskit: QSVM, Qiskit 0.19.6 documentation. `https://qiskit.org/documentation/stubs/qiskit.aqua.algorithms.QSVM.html`.

Rapach, D. E., Strauss, J. K., and Zhou, G. (2013). International stock return predictability: What is the role of the united states? *The Journal of Finance*, 68(4):1633–1662.

Rebentrost, P., Gupt, B., and Bromley, T. R. (2018). Quantum computational finance: Monte Carlo pricing of financial derivatives. *Physical Review A*, 98(2):022321–(1–15).

Rebentrost, P. and Lloyd, S. (2018). Quantum computational finance: quantum algorithm for portfolio optimization. *arXiv:1811.03975*, pages 1–18.

Rebentrost, P., Mohseni, M., and Lloyd, S. (2014). Quantum support vector machine for big data classification. *Physical Review Letters*, 113(13):130503–(1–5).

Richardson, A., van Florenstein Mulder, T., and Vehbi, T. (2021). Nowcasting GDP using machine-learning algorithms: A real-time assessment. *International Journal of Forecasting*, 37(2):941–948.

Sarma, A., Chatterjee, R., Gili, K., and Yu, T. (2019). Quantum unsupervised and supervised learning on superconducting processors. *arXiv:1909.04226*, pages 1–6.

Schuld, M. and Killoran, N. (2019). Quantum machine learning in feature hilbert spaces. *Physical Review Letters*, 122(4):040504–(1–6).

Schuld, M. and Petruccione, F. (2018). Supervised learning with quantum computers. *Springer*.

Sivak, V. V., Eickbusch, A., Royer, B., Singh, S., Tsioutsios, I., Ganjam, S., Miano, A., Brock, B. L., Ding, A. Z., Frunzio, L., Girvin, S. M., Schoelkopf, R. J., and Devoret, M. H. (2023). Real-time quantum error correction beyond break-even. *Nature*, 616:50–55.

Smith, M., Yourish, K., Almukhtar, S., Collins, K., Ivory, D., and Harmon, A. (2020). Coronavirus (covid-19) data in the united states. https://github.com/nytimes/covid-19-data. GitHub.

Smith, S. C., Brown, B. J., and Bartlett, S. D. (2023). Local predecoder to reduce the bandwidth and latency of quantum error correction. *Physical Review Applied*, 19:034050–(1–16).

Stamatopoulos, N., Egger, D. J., Sun, Y., Zoufal, C., Iten, R., Shen, N., and Woerner, S. (2020). Option pricing using quantum computers. *Quantum*, 4:291–311.

Timmermann, A. and Granger, C. W. (2004). Efficient market hypothesis and forecasting. *International Journal of Forecasting*, 20(1):15–27.

Tindall, J., Fishman, M., Stoudenmire, M., and Sels, D. (2023). Efficient tensor network simulation of ibm's eagle kicked ising experiment. *arXiv:2306.14887*, pages 1–18.

Wilkens, S. and Moorhouse, J. (2023). Quantum computing for financial risk measurement. *Quantum Information Processing*, 22(51):1–38.

Zhong, H.-S., Deng, Y.-H., Qin, J., Wang, H., Chen, M.-C., Peng, L.-C., Luo, Y.-H., Wu, D., Gong, S.-Q., Su, H., Hu, Y., Hu, P., Yang, X.-Y., Zhang, W.-J., Li, H., Li, Y., Jiang, X., Gan, L., Yang, G., You, L., Wang, Z., Li, L., Liu, N.-L., Renema, J. J., Lu, C.-Y., and Pan, J.-W. (2021). Phase-programmable gaussian boson sampling using stimulated squeezed light. *Physical Review Letters*, 127(18):180502–(1–9).

Zhou, Y., Stoudenmire, E. M., and Waintal, X. (2020). What limits the simulation of quantum computers? *Physical review X*, 10(041038):1–15.

Zhu, B., Ye, S., Wang, P., Chevallier, J., and Wei, Y.-M. (2022a). Forecasting carbon price using a multi-objective least squares support vector machine with mixture kernels. *Journal of Forecasting*, 41(1):100–117.

Zhu, B., Zhang, J., Wan, C., Chevallier, J., and Wang, P. (2023). An evolutionary cost-sensitive support vector machine for carbon price trend forecasting. *Journal of Forecasting*, 42(4):741–755.

Zhu, Q., Cao, S., Chen, F., Chen, M.-C., Chen, X., Chung, T.-H., Deng, H., Du, Y., Fan, D., Gong, M., Guo, C., Guo, C., Guo, S., Han, L., Hong, L., Huang, H.-L., Huo, Y.-H., Li, L., Li, N., Li, S., Li, Y., Liang, F., Lin, C., Lin, J., Qian, H., Qiao, D., Rong, H., Su, H., Sun, L., Wang, L., Wang, S., Wu, D., Wu, Y., Xu, Y., Yan, K., Yang, W., Yang, Y., Ye, Y., Yin, J., Ying, C., Yu, J., Zha, C., Zhang, C., Zhang, H., Zhang, K., Zhang, Y., Zhao, H., Zhao, Y., Zhou, L., Lu, C.-Y., Peng, C.-Z., Zhu, X., and Pan, J.-W. (2022b). Quantum computational advantage via 60-qubit 24-cycle random circuit sampling. *Science Bulletin*, 67(3):240–245.

8 List of Tables

Table 1: Future forecasts' metrics

Notes. This table displays the means of the future forecasts' balanced accuracies, F1-Scores, Precisions & Recalls.

Algo	C	Gamma	TrainSize	TestSize	nbrObs	bAcc (%)	F1S (%)	Prec (%)	Recall (%)
SVC_rbf	0.01	0.01	80	20	113	49.681	70.154	56.336	92.952
SVC_rbf	0.01	D^{-1}	80	20	113	49.681	70.154	56.336	92.952
SVC_rbf	0.01	1	80	20	113	49.681	70.154	56.336	92.952
SVC_rbf	0.01	100	80	20	113	**49.724**	**70.283**	56.359	93.344
SVC_rbf	0.01	$(D.Var(X))^{-1}$	80	20	113	49.622	70.242	56.306	93.344
SVC_rbf	1	0.01	80	20	113	49.681	70.154	56.336	92.952
SVC_rbf	1	D^{-1}	80	20	113	49.681	70.154	56.336	92.952
SVC_rbf	1	1	80	20	113	49.681	**70.154**	56.336	92.952
SVC_rbf	1	100	80	20	113	50.825	69.169	56.969	88.019
SVC_rbf	1	$(D.Var(X))^{-1}$	80	20	113	**54.633**	67.747	59.561	78.543
SVC_rbf	100	0.01	80	20	113	49.681	70.154	56.336	92.952
SVC_rbf	100	D^{-1}	80	20	113	50.999	**70.895**	57.034	93.657
SVC_rbf	100	1	80	20	113	51.377	69.209	57.290	87.392
SVC_rbf	100	100	80	20	113	**52.531**	65.642	58.207	75.255
SVC_rbf	100	$(D.Var(X))^{-1}$	80	20	113	50.592	57.476	57.011	57.948
SVC_lin	0.01		80	20	113	49.681	70.154	56.336	92.952
SVC_lin	1		80	20	113	49.681	**70.154**	56.336	92.952
SVC_lin	100		80	20	113	**50.486**	70.009	56.767	91.308
QSVC	0.01		80	20	113	49.352	**70.354**	56.168	94.127
QSVC	1		80	20	113	49.502	68.306	56.225	87.001
QSVC	100		80	20	113	**50.804**	58.837	57.164	60.611

Table 2: Present forecasts' metrics

Notes. This table displays the means of the present forecasts' balanced accuracies, F1-Scores, Precisions & Recalls.

Algo	C	Gamma	TrainSize	TestSize	nbrObs	bAcc (%)	F1S (%)	Prec (%)	Recall (%)
SVC_rbf	0.01	0.01	80	20	113	49.889	67.308	54.495	87.997
SVC_rbf	0.01	D^{-1}	80	20	113	49.889	67.308	54.495	87.997
SVC_rbf	0.01	1	80	20	113	49.930	67.349	54.518	88.078
SVC_rbf	0.01	100	80	20	113	**50.498**	**67.922**	54.835	89.213
SVC_rbf	0.01	$(D.Var(X))^{-1}$	80	20	113	50.457	67.881	54.813	89.132
SVC_rbf	1	0.01	80	20	113	49.889	67.308	54.495	87.997
SVC_rbf	1	D^{-1}	80	20	113	49.889	67.308	54.495	87.997
SVC_rbf	1	1	80	20	113	49.930	67.349	54.518	88.078
SVC_rbf	1	100	80	20	113	90.102	91.539	88.914	94.323
SVC_rbf	1	$(D.Var(X))^{-1}$	80	20	113	**91.393**	**92.246**	91.874	92.620
SVC_rbf	100	0.01	80	20	113	49.889	67.308	54.495	87.997
SVC_rbf	100	D^{-1}	80	20	113	88.276	90.407	85.798	95.539
SVC_rbf	100	1	80	20	113	89.777	91.350	88.275	94.647
SVC_rbf	100	100	80	20	113	**91.839**	**92.502**	92.957	92.052
SVC_rbf	100	$(D.Var(X))^{-1}$	80	20	113	90.955	91.443	93.255	89.700
SVC_lin	0.01		80	20	113	49.889	67.308	54.495	87.997
SVC_lin	1		80	20	113	49.889	67.308	54.495	87.997
SVC_lin	100		80	20	113	**88.276**	**90.407**	85.798	95.539
QSVC	0.01		80	20	113	50.416	67.841	54.790	89.051
QSVC	1		80	20	113	82.336	**86.635**	78.272	96.999
QSVC	100		80	20	113	**83.726**	85.805	83.719	87.997

9 List of Figures

Figure 1: Second-order Pauli-Z evolution circuit

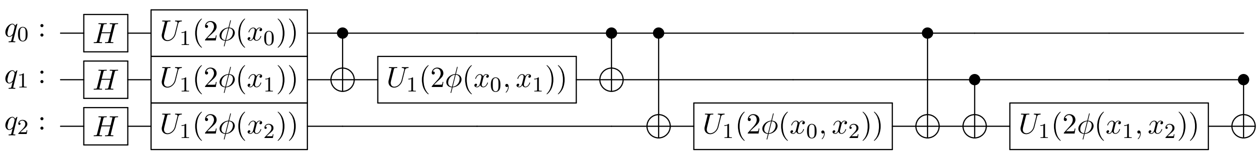

Notes. This chart displays a second-order Pauli-Z evolution circuit for 3 qubits, one repetition and full entanglement. Where ϕ is a classical non-linear function with $\phi(x_i) = x_i$ and $\phi(x_i, x_j) = (\pi - x_i)(\pi - x_j)$.

Figure 2: TSX index weekly movement direction

Notes. These graphs display the up (black +) and down (grey o) weekly movement directions for the S&P/TSX composite index over 551 dates with regards to the one-week lag returns of the S&P/TSX composite and the S&P 500 indexes, and the USDCAD exchange rate.

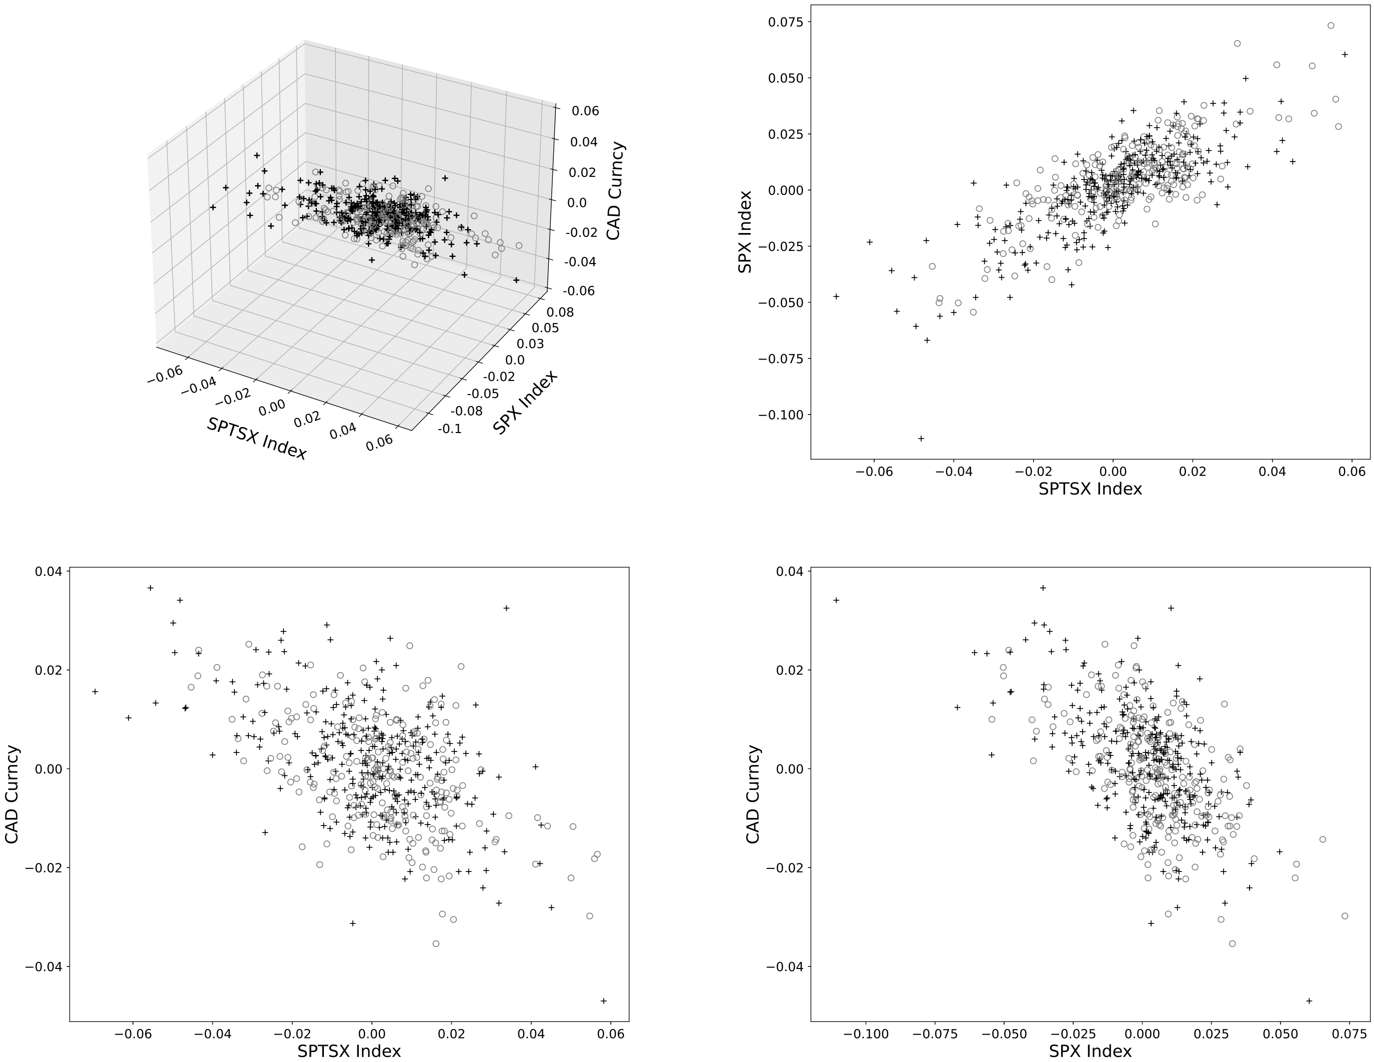

Figure 3: SX5E index weekly movement direction

Notes. These graphs display the up (black +) and down (grey o) weekly movement directions for the Euro Stoxx 50 index over 551 dates with regards to the present returns of the DAX and CAC indexes.

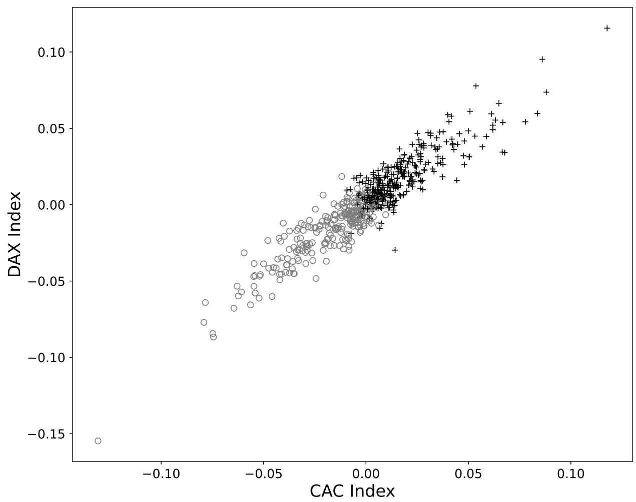

Figure 4: BoC O/N lending rate.

Notes. This graph displays the overnight lending rate set by the Bank of Canada (BoC) on eight fixed dates each year, from December 31, 1997 until October 1, 2021.