Quantum Machine Learning for Financial Applications

Fernando Alonso, Álvaro Leitao, and Carlos Vázquez

Centro de Investigación CITIC, Universidade da Coruña, 15071 A Coruña
M2NICA, Centro de Investigación CITIC, Universidade da Coruña, 15071 A Coruña
M2NICA, Faculty of Computer Science, Universidade da Coruña, 15071 A Coruña
Correspondence: f.alonsof@udc.es

DOI: *https://doi.org/10.17979/spu.23.c19*

Abstract: The rise of quantum technologies has led to growing interest in exploring the use of quantum circuits alongside ML models—also referred to as Quantum Machine Learning (QML)—to enhance the development of new applications, such as through Parametrized Quantum Circuits (PQCs). Among the applications enabled by the use of PQCs in the financial world, some of the most relevant include the computation of risk metrics, as well as the option pricing of financial derivatives. This work aims to illustrate the study of the aforementioned applications, where the main idea consists in approximating, through the PQC, the underlying distribution of the assets and the payoff function, and from there, computing risk metrics and derivative prices.

1 Introduction

A fundamental task in quantitative finance is **derivatives pricing**. This consists in determining the price of the derivative at any date prior to its maturity, which requires the future and uncertain dynamics of the underlying asset. Such dynamics are typically modeled by stochastic differential equations (SDEs). Given a financial option or derivative, let S_t be the **price of the underlying asset** at time $t \in [0, T]$, where T is the maturity date, and let the **price of the derivative** be $v_t = V(t, S_t)$, for some function $V : [0, T] \times \mathbb{R} \to \mathbb{R}$. Given a **payoff function** defined as $v_T = h(T, S_T)$, the option price can be formulated in terms of a conditional expectation as:

$$v(x, t_0) = e^{-r(T-t_0)}\, \mathbb{E}^{\mathbb{Q}}[h(T, y)|x] = e^{-r(T-t_0)} \int_{\mathbb{R}} h(T, y) p(y|x)\, dy, \tag{10.1}$$

where x and y are the state variables at times t_0 and T, respectively, and $p(y|x)$ is the probability density of y conditioned on x.

In view of the above expression, the valuation of these financial derivatives primarily requires the computation of the expectation involved. One of the best-known integration methods for solving the pricing problem is the Monte Carlo method. Since this method can be computationally demanding for certain types of integrals, in recent years the advantages offered by quantum computing have been exploited to develop the *Quantum Monte Carlo Integration* (QMCI) method (Gómez et al. (2022); Montanaro (2015)), which presents some challenges in today's hardware. The quantum methods proposed in this work draw some inspiration from the QMCI variants, particularly from those that integrate QML-based techniques to approximate Fourier series, with the goal of leveraging the expressive power and generalization capabilities that these models offer.

1.1 PQCs as universal aproximators

A very common approach within the quantum-classical framework of QML consists in using trainable quantum circuits as models, in a similar way to neural networks. In this approach, the circuit is measured multiple times to estimate the expected value of an observable and the parameters are generally adjusted according to a criterion based on this output, leading to the implementation of a function $f_\theta(x)$. This approach is referred in the literature as PQCs. In this context, the fundamental idea of our approach is that, given a quantum circuit, $U(x, \boldsymbol{\theta})$, which depends on the input (x) and on a set of parameters $(\boldsymbol{\theta})$ and an observable, M; we can define an univariate quantum model that can be written as a Fourier series, i.e.,

$$f_\theta(x) = \langle 0|U^\dagger(x, \boldsymbol{\theta}) M U(x, \boldsymbol{\theta})|0\rangle = \sum_{\omega \in \Omega} c_\omega(\theta) e^{i\omega x}.$$

The quantum circuit is constructed with L layers, each consisting of a data-encoding block $S(x)$ and a trainable block $W(\boldsymbol{\theta})$, as illustrated in *Figure* 1.

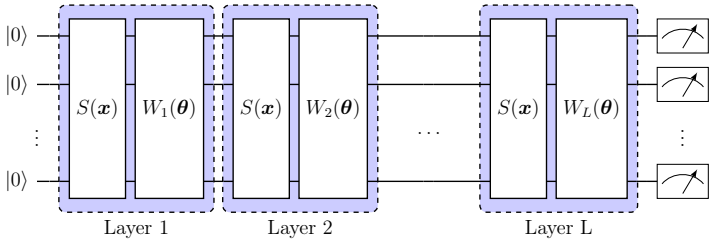

Figure 1: Scheme of a quantum circuit composed of L layers, from Manzano-Herrero (2024).

Several references in the literature, for example in Schuld et al. (2021), explore two key aspects of PQCs: the **universality** and the **expressivity**, that is, the ability of PQCs to approximate any function belonging to a given function space defined over a prescribed domain, up to arbitrary precision with respect to a specific norm. Furthermore, Manzano et al. (2025) and *Chapter* 5 of Manzano-Herrero (2024) prove that PQCs can approximate arbitrarily well the space of continuous functions, the space of p-integrable functions and the Sobolev space H^k.

2 Pricing method formulation

As mentioned before, in quantitative finance, efficient numerical methods are required to value complex contracts and calibrate financial models. In this work, two classical-quantum methods have been developed, employing trigonometric Fourier series.

2.1 PDF approximation formulation

The starting point lies in the calculation of the pricing formula under the risk-neutral measure given by (10.1). In practice, in finance is common to work with probability densities whose tails tend to vanish, so it can be assumed that there exists an interval $[a, b] \subset \mathbb{R}$ such that the integral can be approximated without significant loss of accuracy, i.e.,[1]

$$v_1(x, t_0) = e^{-r(T-t_0)} \int_a^b h(T, y) p(y|x)\, dy. \tag{10.2}$$

[1] Subscripts are used to denote successive numerical approximations.

Since $p(y|x)$ is usually not known, it can be replaced by its trigonometric expansion

$$p(y|x) \approx \frac{A_0}{2} + \sum_{k=1}^{N} \left(A_k \cos\left(2\pi k \frac{(y-a)}{b-a}\right) + B_k \sin\left(2\pi k \frac{(y-a)}{b-a}\right)\right), \tag{10.3}$$

where

$$A_k = \frac{2}{b-a}\int_a^b p(y|x)\cos\left(2\pi k \frac{(y-a)}{b-a}\right)\, dy \text{ and } B_k = \frac{2}{b-a}\int_a^b p(y|x)\sin\left(2\pi k \frac{(y-a)}{b-a}\right)\, dy.$$

It should be remarked that regarding the underlying asset price process, we can use the same arguments as in Fang and Oosterlee (2009) to ensure that, due to the conditions required for the existence of the Fourier series, it is possible to truncate the number of terms in the series while controlling the accuracy.

By substituting equation (10.3) into equation (10.2), exchanging the summation and the integral and introduce the definition

$$C_k := \frac{2}{b-a}\int_a^b h(T,y)\cos\left(2\pi k \frac{(y-a)}{b-a}\right)\, dy, \text{ and } D_k := \frac{2}{b-a}\int_a^b h(T,y)\sin\left(2\pi k \frac{(y-a)}{b-a}\right)\, dy,$$

we obtain

$$v_2(x,t_0) \approx \frac{1}{2}(b-a)e^{-r\Delta t}\left(\frac{A_0 \cdot C_0}{2} + \sum_{k=1}^{N}\left(A_k \cdot C_k + B_k \cdot D_k\right)\right). \tag{10.4}$$

Note that C_k and D_k are the sine and cosine series coefficients of $h(T,y)$. Thus, the integral of the product of two real functions, $p(y|x)$ and $h(T,y)$ has been transformed into the sum of the product of their coefficients of the trigonometric expansion.

3 Methodology

The proposed method is based on the use of a QML model built on PQCs, with the aim of approximating the PDF through their Fourier series and extracting the coefficients needed to compute the price of a financial derivative. To carry out this task, two different approaches have been designed. In the first one, learning is supervised, using datasets that contain both inputs and labeled outputs. The second approach is more realistic, since in derivative pricing one rarely has access to the exact probability distribution of the underlying asset price. Therefore, self-supervised learning is employed, providing the model with a sufficiently representative set of asset samples from which the model must infer the implicit distribution. It is worth remarking that we will work in the frame of *Quantum Differential Machine Learning*,

3.1 Method I: Supervised Learning for PDF approximation

First, we consider a labeled dataset associated to the probability distribution, defined over a truncation interval that ensures the Fourier series approximation is sufficiently accurate. Then, we train the PQC in order to approximate the underlying probability density function in $[-2\pi, 2\pi]$ rescaling the data to $[-\pi, \pi]$. This ensures that the resulting series is smoother and does not exhibit Gibbs phenomena, since it is allowed freedom of approximation outside $[-\pi, \pi]$, which can be observed in *Figure* 2.

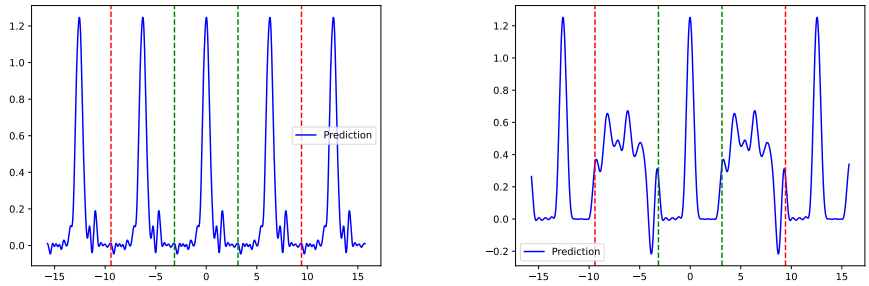

(a) (b)

Figure 2: Approximation of $\log(S_t/K)$ without freedom (a) and with freedom (b).

The Fourier coefficients are extracted by the **Discrete Fourier Transform** (DFT), applied to direct evaluations of the trained circuit in $[-\pi, \pi]$. Once this is done, the recovered coefficients $c_k \in \mathbb{C}$ correspond to the exponential series of order N, which are transformed in the desired trigonometric coefficients A_k and B_k via the Euler identities. The Fourier coefficients associated with the payoff, C_k and D_k, are obtained analytically. The final price is computed by (10.4). The empirical risk function is

$$R_{h^1}^S(f^*, f) = \frac{1}{n}\sum_{i=0}^{n-1}\left(f^*(x_i) - f(x_i)\right)^2 + \sum_{j=0}^{N-1}\frac{1}{n}\sum_{i=0}^{n-1}\left(\frac{\partial f^*}{\partial x_j}(x_i) - \frac{\partial f}{\partial x_j}(x_i)\right)^2.$$

3.2 Method II: Self-supervised Learning for PDF approximation

As mentioned before, in this approach **self-supervised** learning will be employed, providing the model with only a representative set of asset samples, so that the labels required to construct the model's cost function are generated internally from that set. Therefore, it is necessary to define a new risk function that uses only the inputs $x \in \mathcal{X}$ and somehow incorporates the derivatives to ensure convergence of the approximation. Following the formulation based on the results of *Chapter* 7 of Manzano-Herrero (2024), where, from various convergence results an adaptation of the classical optimization problem is proposed; we consider the empirical risk function:

$$R_{l^2,J^2}^{\mathcal{T}}(F) = R_{\mathcal{T},l^2}^{\mathcal{X}}(F) + R_{l^2}^{\mathcal{T}}(f) = \frac{1}{n}\sum_{i=0}^{n-1}\left(F_{\text{emp}}^*(x_i) - F(x_i)\right)^2 - \frac{2}{n}\sum_{i=0}^{n-1} f(x_i) + Q(f^2).$$

Onwards, the way to proceed is the same as in Method I, taking into account that it is necessary to differentiate the series to obtain the PDF. This may involve an additional error in the results obtained.

4 Results

In the following section, we present and analyze the obtained results.

4.1 Experiment setting

In the experiments, we consider European vanilla options and assume that the underlying asset follows Black–Scholes modelc, with the following fixed parameters:

$$S_0 = 100, \quad r = 0.1, \quad q = 0, \quad T = 1, \quad \sigma = 0.25.$$

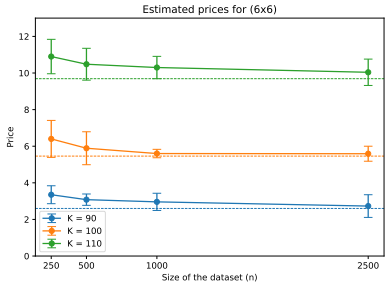

(a)

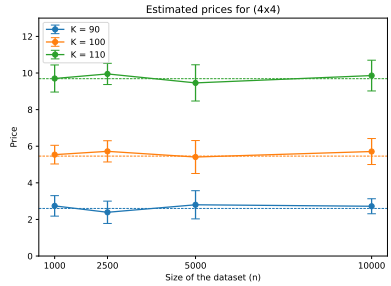

(b)

Figure 3: Comparison of Method I with 6×6 scheme (a) and Method II with 4×4 scheme (b).

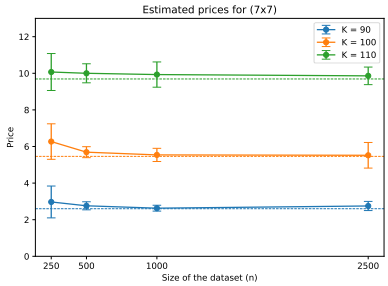

(a)

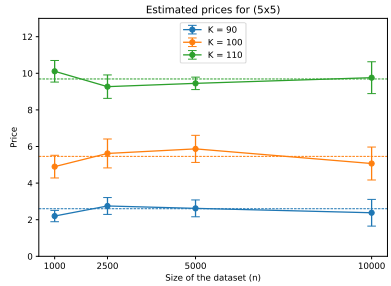

(b)

Figure 4: Comparison of Method I with 7×7 scheme (a) and Method II with 5×5 scheme (b).

Three different strike prices have been selected, $K \in \{90, 100, 110\}$, with the aim of analyzing the model's behavior for different option contract configurations. All experiments are conducted using the same quantum *ansatz* as in Manzano-Herrero (2024).

4.2 Impact of the number of coefficients

Since the functions to be approximated in each method are different, it is necessary to use different dimensions in the PQC schemes to ensure good results. This is because, in certain cases, adding an excessive number of terms in the approximation of a function can introduce additional contributions to oscillations and instabilities in the model. In the specific case of the CDF, which exhibits sharp jumps at its boundaries due to periodic extension, these oscillations may become more pronounced.

Nevertheless, when using higher dimensions both methods display stable behavior, showing a clear trend of convergence towards the true derivative value as the expressive capacity of the circuit increases. This can be clearly seen from *Figures* 3 to 5, reinforcing the idea that introducing flexibility in the Fourier series approximation is essential to achieve accurate and reliable estimates.

It is worth noting that, according to the results, the best performance is obtained with an intermediate scheme as shown in *Figure* 6. This shows how in this type of problems there is often a *trade-off* between model complexity and the specific characteristics of the problem under consideration.

Despite these limitations, the results obtained so far are clearly satisfactory and demonstrate the potential of the approach.

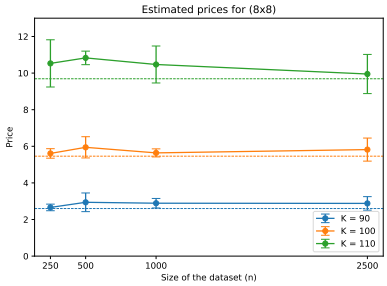

(a)

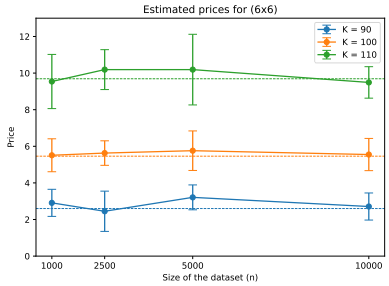

(b)

Figure 5: Comparison of Method I with 8×8 scheme (a) and Method II with 6×6 scheme (b).

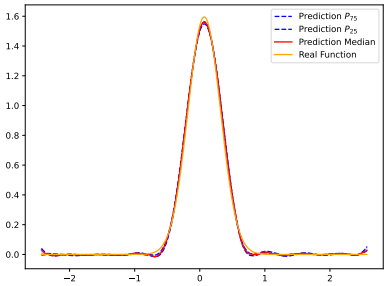

(a)

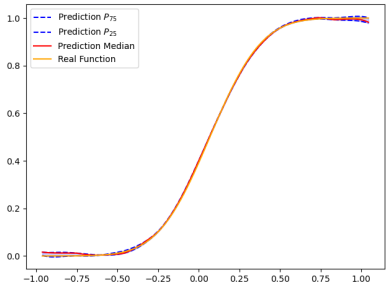

(b)

Figure 6: Output of the PQC for Method I with 7×7 scheme (a) and for Method II with 5×5 scheme (b).

4.3 Impact of the data size

As the number of observations increases, the estimations obtained through approximation models should reflect more accurately the statistical patterns of the original data, thus resulting in a better generalization and a reduction of the error.

From *Figures* 3 to 5, a clear convergence towards the exact value can be observed. However, some outliers appear, which can be attributed to an inherent error of the training process, particularly to the randomness in the initialization of the quantum model weights, which introduces fluctuations that are not always corrected during optimization.

5 Conclusions

A classical–quantum method for the pricing of financial derivatives has been proposed and formulated, based on the outputs of QML models built with PQCs. Within this framework, two different training approaches have been developed, allowing the robustness and versatility of the model to be assessed in various practical scenarios. The results show that, for both methods, using larger sample sizes and more complex ansatz configurations leads to a clear tendency toward convergence to the true values. Finally, it has been shown that Method II is particularly promising for real-world situations in which only sample data of the underlying asset are available, thus representing an innovative alternative to traditional option pricing techniques.

Bibliography

F. Fang and C. W. Oosterlee. A novel pricing method for european options based on fourier-cosine series expansions. *SIAM Journal on Scientific Computing*, 31(2):826–848, 2009.

A. Gómez, A. Leitao, A. Manzano, M. Nogueiras, G. Ordóñez, and C. Vázquez. A survey on quantum computational finance for derivatives pricing and var. *Archives of Computational Methods in Engineering*, 03 2022.

A. Manzano, D. Dechant, J. Tura, and V. Dunjko. Approximation and generalization capacities of parametrized quantum circuits for functions in sobolev spaces. *Quantum*, 9, 2025.

A. P. Manzano-Herrero. *Contributions to the Pricing of Financial Derivatives Contracts in Commodity Markets and the Use of Quantum Computing in Finance*. Doctoral thesis, Universidade da Coruña (UDC), 2024.

A. Montanaro. Quantum speedup of monte carlo methods. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 471(2181), 2015.

M. Schuld, R. Sweke, and J. J. Meyer. Effect of data encoding on the expressive eower of variational quantum-machine-learning models. *Physical Review A*, 103(3), 2021.